

Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate

Stephen Ansolabehere
Harvard University

Eitan Hersh
Institution for Social and Policy Studies, Yale University, New Haven, CT 06520-8209
e-mail: eitan.hersh@yale.edu (corresponding author)

Edited by R. Michael Alvarez

Social scientists rely on surveys to explain political behavior. From consistent overreporting of voter turnout, it is evident that responses on survey items may be unreliable and lead scholars to incorrectly estimate the correlates of participation. Leveraging developments in technology and improvements in public records, we conduct the first-ever fifty-state vote validation. We parse overreporting due to response bias from overreporting due to inaccurate respondents. We find that nonvoters who are politically engaged and equipped with politically relevant resources consistently misreport that they voted. This finding cannot be explained by faulty registration records, which we measure with new indicators of election administration quality. Respondents are found to misreport only on survey items associated with socially desirable outcomes, which we find by validating items beyond voting, like race and party. We show that studies of representation and participation based on survey reports dramatically misestimate the differences between voters and nonvoters.

1 Introduction

Survey research provides the foundation for the scientific understanding of voting. Yet, there is a nagging doubt about the veracity of research on political participation because the rate at which people report voting in surveys greatly exceeds the rate at which they actually vote. For example, 78% of respondents to the 2008 National Election Study (NES) reported voting in the presidential election, compared with the estimated 57% who actually voted (McDonald 2011)—a twenty-one-point deviation of the survey from actuality.¹ That difference is comparable to the total effect of variables such as age and education on voting rates, and such bias is almost always ignored in analyses that project the implications of correlations from surveys onto the electorate, such as studies of the differences in the electorate “were all people to vote.”

Concerns over survey validity and the correct interpretation of participation models have made vote misreporting a continuous topic of scholarly research, as evidenced by recent work by Katz and Katz (2010), Campbell (2010), and Deufel and Kedar (2010). To correct for misreporting, there have been a number of attempts in the past to validate survey responses with data collected from government sources, though no national survey has been validated in over twenty years. Survey validation has been met with both substantive and methodological critiques, such as that validation techniques are error-prone and prohibitively costly and that the misestimation of turnout is primarily a function of sample selection bias rather than item-specific misreporting (e.g., Berent, Krosnick, and Lupia 2011).

A new era characterized by high-quality public registration records, national commercial voter lists, and new technologies for big data management creates an opportunity to revisit survey

Authors' note: We thank the editors and anonymous reviewers for their helpful feedback. For materials to replicate the statistical analysis in this article, see Ansolabehere and Hersh (2012).

¹In the NES vote validation in the 1970s and 1980s, about 20%–30% of nonvoters claimed to have voted (Traugott, Traugott, and Presser 1992; Belli, Traugott, and Beckmann 2001, 483).

validation. The new tools that facilitate a less costly and more reliable match between survey responses and public records provide a clearer picture of the electorate than was ever before possible. This article presents results from the first validation of a national voter survey since the NES discontinued its vote validation program in 1990. It is the first-ever validation of a political survey in all fifty states. The study was conducted by partnering with a commercial data vendor to gain access to all states' voter files, exploit quality controls and matching technology, and compare information on commercially available records with the survey reports. We validate not just voting reports (which were the focus of the NES validation), but also whether respondents are registered or not, the party with which respondents are registered, respondents' races, and the method by which they voted. Validation of these additional pieces of information provides important clues about the nature of validation and misreporting in surveys.

Several key findings emerge from this endeavor. First, we find that standard predictors of participation, like demographics and measures of partisanship and political engagement, explain a third to a half as much about voting participation as one would find from analyzing behavior reported by survey respondents. Second, in the Web-based survey we validate, we find that most of the overreporting of turnout is attributable to misreporting rather than to sample selection bias.² This is surprising, inasmuch as the increasing use of Internet surveys raises the concern that the Internet mode is particularly susceptible to issues of sampling frame. Third, we find that respondents regularly misreport their voting history and registration status, but almost never misreport other items on their public record, such as their race, party, and how they voted (e.g., by mail, in person). Whatever process leads to misreporting on surveys, it does not affect all survey items in the same way. Fourth, we employ individual-level, county-level, and state-level measures of data quality to test whether misreporting is a function of the official records rather than respondent recall. We find that data quality is hardly at all predictive of misreporting.

After detailing how recent technological advancements and a partnership with a private data vendor allow us to validate surveys in a new way, we compare validated voting statistics to reported voting statistics. Following the work of Silver, Anderson, and Abramson (1986), Belli, Traugott, and Beckmann (2001), Fullerton, Dixon, and Borch (2007), and others, we examine the correlates of survey misreporting. We compare misreporting in the 2008 Cooperative Congressional Election Study (CCES) to misreporting in the validated NES from the 1980, 1984, and 1988 Presidential elections, and find that in spite of the time gap between the validation studies as well as the differences in survey modes and validation procedures, there is a high level of consistency in the types of respondents who misreport.

In demonstrating a level of consistency with previous validations and finding stable patterns of misreporting that cannot be attributed to sample selection bias or faulty government records, this analysis reaches different conclusions than a recent working paper of the NES, which asserts that self-reported turnout is no less accurate than validated turnout and therefore that survey researchers should continue to rely on reported election behavior (Berent, Krosnick, and Lupia 2011). It also combats prominent defenses of using reported vote rather than validated vote in studies of political behavior (e.g., Verba, Schlozman, and Brady 1995, Appendix A). We find the evidence from the 2008 CCES validation convincing that electronic validation of survey responses with commercial records provides a far more accurate picture of the American electorate than survey responses alone.

The new validation method we employ addresses the problems that have plagued past attempts at validating political surveys. As such, the chief contribution of the article is to address concerns with survey validation, describe and test the ways that new data and technology allow for a more reliable matching procedure, and show how electronic validation improves our understanding of the electorate. Apart from its methodological contributions, this article also emphasizes the limitations of "resources"-based theoretical models of participation (e.g., Rosenstone and Hansen 1993; Verba, Schlozman, and Brady 1995). Such models do not take us very far in explaining who votes and who abstains; rather, they perform the dubious function of predicting the types

²In the NES studies, misreporting and sample selection contribute to overreporting in equal parts.

of people who think of themselves as voters when responding to surveys. Demonstrating how little resources like education and income correlate with voting, this research calls for more theory-building if we are to successfully capture the true causes of voting. Since misreporters look like voters, reported vote models exaggerate the demographic and attitudinal differences between voters and non-voters. This finding is consistent with research by Cassel (2003) and Bernstein, Chaha, and Montjoy (2001) (see also Highton and Wolfinger 2001 and Citrin, Schickler, and Sides 2003).

To summarize, the contributions of this essay are as follows. The validation project described herein is the first-ever fifty-state vote validation. It is the first national validation of an Internet survey, a mode of survey analysis on the rise. It is the first political validation project that considers not only voting and registration, but also the accuracy of respondents' claims about their party affiliation, racial identity, and method of voting. It is the first political validation project that uses individual-, county-, and state-level measures of registration list quality to distinguish misreporting attributable to poor records from misreporting attributable to respondents inaccurately recalling their behavior. Our efforts here sort out misreporting from sample selection as contributing factors to the misestimation of election participation, sort out registration quality from respondent recall as contributing factors to misreporting, show that a consistent type of respondent misreports across survey modes and election years, find that responses on nonsocially desirable items are highly reliable, and identify the correlates of true participation as compared with reported participation.

2 Why Do Survey Respondents Misreport Behavior?

Scholars have articulated five main hypotheses for why vote overreporting shows up in every study of political attitudes and behaviors. First, the aggregate overreporting witnessed in surveys may not result from respondents inaccurately recalling their participation but rather may be an artifact of sample selection bias. Surveys are voluntary, and if volunteers for political surveys come disproportionately from the ranks of the politically engaged, then this could result in inflated rates of participation. Along these lines, Burden (2000) shows that NES respondents who were harder to recruit for participation were less likely to vote than respondents who immediately consented to participate.³ There is little doubt that sample selection contributes at least in part to the overreporting problem, and here we attempt to measure how big of a part it is. However, the patterns of inconsistencies between reported participation and validated participation as identified in the NES vote validation efforts suggest that sample selection is not the only phenomenon leading to overreporting.

Second, survey respondents may forget whether they participated in a recent election or not. Politics is not central to most people's day-to-day lives, and they might just fail to remember whether they voted in a specific election. In support of the memory hypothesis, Belli et al. (1999) note that as time elapses from Election Day to the NES interviews, misreporting seems to occur at a higher rate (though Duff et al. 2007 do not find such an effect). In an experiment, Belli et al. (2006) show that a longer-form question about voting that focuses on memory and encourages face-saving answers reduces the rate of reported turnout.

Memory may play some role in misreporting; however, a couple of facts cut against a straightforward memory hypothesis. Virtually no one who is validated as having voted recalls in any survey that they did not vote. If memory alone was the explanation for misreporting, we would expect an equal number of misremembering voters and misremembering nonvoters, but it is only the nonvoters who misremember. Second, in the case of the 2008 validated CCES study, 98% of the post-election interviews took place within two weeks of the Presidential election. Even for those whose interest in politics is low, the Presidential election was a major world event. It is somewhat hard to imagine a large percentage of people really failing to remember if they voted in a high-salience election that took place not longer than two weeks prior. And as we will see, the

³For a discussion of sample selection, see the exchange following Burden's (2000) article (Burden 2003; Martinez 2003; McDonald 2003). Martinez focuses on the effect of panel studies on overreporting wherein attrition is most common among those least likely to participate.

respondents who do misreport are not the ones who are disengaged with politics; on the contrary, misreporters claim to be very interested in politics.

A third hypothesis emanating from the NES validation studies involves an in-person interview effect (Katosh and Traugott 1981). Perhaps when asked face-to-face about a socially desirable activity like voting, respondents tend to lie. However, Silver, Abramson, and Anderson (1986) note similar rates of misreporting on telephone and mail surveys as in face-to-face surveys. With Web surveys, as in mail surveys, the impersonal survey experience may make it easier for a respondent to admit to an undesirable behavior but may also make lying a cognitively easier thing to do, as an experiment by Denscombe (2006) suggests. Malhotra and Krosnick (2007) posit that social desirability bias may be lower in an in-person interview because of the “trust and rapport” built between people meeting in person. We show here that misreporting is quite common in Internet-based surveys—more common, in fact, than in the NES in-person studies, which might be due to an over sample of politically knowledgeable respondents. At the margins, misreporting levels may rise or fall by survey mode, but it does not appear that any survey mode in itself explains misreporting.

A fourth hypothesis for vote-overreporting is that the inconsistencies between participation as reported on surveys and participation as indicated in government records are an artifact of poor record-keeping rather than false reporting by respondents. This hypothesis is most clearly articulated in recent work by Berent, Krosnick, and Lupia (2011), who argue that record-keeping errors and poor government data make the matching of survey respondents to voter files unreliable. Record-keeping problems were particularly concerning prior to federal legislation that spurred the digitization and centralization of voting records and created uniform requirements for the purging of obsolete records. However, even during this period—the period in which all NES studies were validated—record-keeping did not seem to be the main culprit of vote misreporting. As Cassel (2004) finds, following up on the work of Presser, Traugott, and Traugott (1990), only 2% of NES respondents who were classified as having misreported were misclassified on account of poor record-keeping. The NES discontinued its validation program not because of concerns about record quality, but rather on account of the cost associated with an in-person validation procedure (Rosenstone and Legee 1994).

While record-keeping was and continues to be a concern in survey validation, several facts cut against the view that voters are generally reporting their turnout behavior truthfully but problems with government records lead to the appearance of widespread misreporting. If poor record-keeping was the main culprit, one would not expect to find consistent patterns across years and survey modes of the same kinds of people misreporting. Nor would one expect to find that only validated nonvoters misreport. Nor would one expect that validated and reported comparisons on survey items like race, party, and voting method would be consistent but only socially desirable survey items would be inconsistent, as we find here. Apart from these findings, we will go further here and utilize county-level measures of registration list quality to show that while a small portion of misreporting can be explained by measures of election administration quality, they do not explain nearly as much as personality traits, such as interest in politics, partisanship, and education. Together, these pieces of evidence refute the claim that misreporting is primarily an artifact of poor record-keeping. As we articulate the process of modern survey matching below, we will return to the work by Berent, Krosnick, and Lupia (2011) and explain why we reach such different conclusions than they do.

The fifth hypothesis for vote overreporting, and the most dominant one, is best summarized by Bernstein, Chaha, and Montjoy (2001, 24): “people who are under the most pressure to vote are the ones most likely to misrepresent their behavior when they fail to do so.” Likewise, Belli, Traugott, and Beckmann (2001) argue that over-reporters are those who are similar to true voters in their attitudes regarding the political process.⁴ This argument is consistent with the finding that over-reporters seem to look similar to validated voters (Sigelman 1982). For example, like validated voters, over-reporters tend to be well-educated, partisan, older, and regular church attendees.

⁴Also see Karp and Brockington (2005) and Harbough (1996).

Across all validation studies, education is the most consistent predictor of overreporting, with well-educated respondents more likely to misreport (Silver, Anderson, and Abramson 1986; Bernstein, Chaha, and Montjoy 2001; Belli, Traugott, and Beckmann 2001; Cassel 2003; Fullerton, Dixon, and Borch 2007).⁵ The social desirability hypothesis is also supported by experimental work: survey questions that provide “socially acceptable excuses for not voting” have reduced overreporting by eight percentage points (Duff et al. 2007).

2.1 *Does Misreporting Bias Studies of Voting?*

Whatever degree to which social desirability and these other phenomena contribute to misreporting, there is a separate question of whether misreporting leads to faulty inferences about political participation. Canonical works of voting have defended the study of reported behavior over validated behavior not merely on the grounds of practicality, but by suggesting that the presence of misreporters actually does not bias results (Wolfinger and Rosenstone 1980; Rosenstone and Hansen 1993; Verba, Scholzman, and Brady 1995). As we argue in detail elsewhere (Ansolabehere and Hersh 2011a), these claims are not quite right. Apart from using standard regression techniques, these major works on political participation all rely extensively on analyses of statistics that Rosenstone and Hansen call “representation ratios” and Verba et al. call “logged representation scales.” These statistics flip the usual conditional relationship of estimating voting given a person’s characteristics, instead studying characteristics given that a person voted. Because validated voters and misreporters look similar on key demographics, these ratio statistics do not fluctuate very much depending on the use of validated or reported behavior. However, the statistics themselves conflate the differences between voters and non-voters with the proportion of voters and non-voters in a sample. Because surveys tend to be disproportionately populated by reported voters, the ratio measures can lead to faulty inferences. Again, we consider these problems in greater detail elsewhere; here, we emphasize that when one is studying the differences between voters and nonvoters, misreporters do indeed bias results.

3 The Commercial Validation Procedure

In the spring of 2010, we entered a partnership with Catalist, LLC. Catalist is a political data vendor that sells detailed registration and microtargeting data to the Democratic Party, unions, and left-of-center interest groups. Catalist and other similar businesses have created national voter registration files in the private market. They regularly collect voter registration data from all states and counties, clean the data, and make the records uniform. They then append hundreds of variables to each record. For example, using the registration addresses, they provide their clients with Census information about the neighborhood in which each voter resides. Using name and address information, they contract with other commercial firms to append data on the consumer habits of each voter. As part of our contract, Catalist matched the 2008 CCES into its national database. Polimetrix, the survey firm that administered the CCES, shared with Catalist the personal identifying information it collected about the respondents. Using this information, including name, address, gender, and birth year, Catalist identified the records of the respondents and sent them to Polimetrix. Polimetrix then de-identified the records and sent them to us.

Using a firm like Catalist solves many of the record-keeping issues identified in the NES validation studies that utilized raw voter registration files from election offices. However, as Berent, Krosnick, and Lupia (2011) note, private companies make their earnings off proprietary models, and so there is potentially less transparency in the validation process when working with an outside

⁵A special case of the social desirability hypothesis is the propensity of African Americans to over-report more than Whites. Contrary to the general pattern of over-reporters and validated voters appearing similar, Blacks may be under heightened social pressure to vote but they also generally vote at lower rates than Whites (Deufel and Kedar 2010). Some have theorized that in the wake of the Civil Rights movements, many Blacks feel duty-bound to vote, and are pressured to vote by their racial cohort (see Belli, Traugott, and Beckmann 2001; Fullerton, Dixon, and Borch 2007; Duff et al. 2007). However, in the 2008 election, in which the first African American major-party candidate ran and won, the rate of vote misreporting among Blacks was not noticeably higher than among Whites, as we show below.

company. For this reason, we provide a substantial amount of detail about the method and quality of Catalist's matching procedure. Three factors give us confidence that Catalist successfully links respondents to their voting record and thus provides a better understanding of the electorate than survey responses alone. These three factors are (1) an understanding of the data-cleansing procedure that precedes matching, which we learned about through over twenty hours of consultation time with Catalist's staff; (2) two independent verifications of Catalist's matching procedure, one by us and one by a third party that hosts an international competition for name-matching technologies; and (3) an investigation of the matched CCES, in which we find strong confirmatory evidence of successful matching. We consider the first two factors now, and the third in our data analysis.

3.1 *Pre-Processing Data: The Key to Matching*

In matching survey respondents to government records, arguably the simplest part of the process is the algorithm that links identifiers between two databases. The more important ingredient to validation is pre-processing and supplementing the government records in order to facilitate a more accurate match. A recent attempt at survey validation by the NES that found survey validation to be difficult and unreliable is notable because this attempt did not take advantage of available tools to pre-process registration data ahead of matching records (Berent, Krosnick, and Lupia 2011). Given a limited budget, this is understandable. The reason name matching is a multi-million-dollar business (Catalist conducted over nine billion matches in 2010 alone) is that it requires large amounts of supplementary data and area-specific expertise, neither of which were part of the NES validation.

For example, one of the challenges of survey validation is purging. As the NES study notes, some jurisdictions take many months to update vote history on registration records, and for this reason it is preferable to wait a year or longer after an election to validate. On the other hand, the longer one waits, the more likely that registration files change because government offices purge records of persistent nonvoters, movers, and the deceased. As with our validation, the NES study validated records from the 2008 general election in the spring of 2010. Unlike our validation, they used a single snapshot of the registration database in 2010, losing all records of voters who were purged since the election. What Catalist does is different. Catalist obtains updated registration records several times a year from each jurisdiction. When Catalist compares a file recently acquired with an older file in its possession, it notes which voters have been purged but retains their voting records. In the 2008 validated CCES, 3% of the matched sample had previously been registered but were since dropped. Another 2% are listed as inactive—a designation that is often the first step to purging voters. The retention of dropped voters is one of the biggest advantages to using a company like Catalist for matching.

Another important pre-processing step is obtaining commercial records from marketing firms. Prominent data aggregation vendors in the United States collect information from credit card companies, consumer surveys, and government sources and maintain lists of U.S. residents that they sell to commercial outlets. Catalist contracts with one such company to improve its matching capability. For instance, suppose a voter is listed in a registration file with a missing birthdate field or with only a year of birth rather than an exact date. Catalist may be able to find the voter's date of birth by matching the person's name and address to a commercial record. In the recent NES validation, a substantial number of records were considered by the researchers to be "probably inaccurate" simply because they were missing birthdate information. However, a perfectly valid registration record might not have complete birthdate information because the jurisdiction does not require it or a handwritten registration application was illegible in this field but was otherwise readable, or for some other reason. In the CCES survey data matched with Catalist records, 17% of matched records used an exact birthdate identified through a commercial vendor rather than a year of birth or blank field as one would use with a raw registration file. Similarly, Catalist runs a simple first-name gender model to impute gender on files that do not include gender in the public record. Twenty-eight percent of our respondents have an imputed gender, which improves the validation

accuracy. Again, these sorts of imputations are common practice in the field of data management, but the NES validation did not use such techniques, which contributed to their low match rate.

Other preprocessing steps that Catalist takes (but the NES validation did not take) include: (1) Catalist de-duplicates records by linking records of the same person listed in a state's voter file more than once. (2) Catalist collects data from all fifty states, so the validation procedure did not require us to go state-to-state or to restrict our analysis to a small subsample of states as the NES has done. (3) Catalist runs all records through the Post Office's National Change of Address (NCOA) Registry to identify movers. Of course, since not every person who moves registers the move with the Post Office, the tracking of movers is imperfect; however, it is a substantial improvement on using a single snapshot from a raw registration file as the basis of a survey validation. (4) Because Catalist opens up its data-processing system to researchers, enabling them to evaluate the raw records it receives from the election offices nationwide, we are able to generate county-level measures of election administration quality and observe how they relate to misreporting. Similarly, when Catalist matches surveys to its voter database, it creates a confidence statistic for each person indicating how sure it is that the match is accurate. We can also condition our analysis of misreporting on this confidence score.

The contrast between the advantages of modern data processing and the techniques used in the recent NES validation is important because the NES led the way in survey validation in the 1970s and 1980s and thus its voice carries weight when it suggests that scholars should continue to rely on reported vote data and that "overestimation of turnout rates by surveys is attributable to factors and processes other than respondent lying" (Berent, Krosnick, and Lupia 2011). Because the NES attempted to validate its survey without the resources and expertise of a professional firm, they failed to take simple steps like tracking purged voters and movers, imputing gender and birthdate values, and de-duplicating records, all steps that are standard practice in the field.^{6,7}

3.2 *Validating the Validator: How Good Are Catalist's Matches?*

Once the government data are pre-processed, Catalist uses an algorithm based on the identifying information transmitted from their clients in order to link records. The algorithm is proprietary, but this is what we can report from meetings with their technical staff. The first stage of matching is a "fishing algorithm." Catalist takes all of the identifying data they receive from a client, extends bounds around each datum, and casts a net to find plausible matches. For example, if they only receive year of birth from a client (as they did for the CCES), they might put a cushion of a few years before and after the listed birth year in the event that there is a slight mismatch. Once the net is cast, they turn to a "filter algorithm" that tries to identify a person within the set of potential matches. The confidence score that Catalist gives its clients is essentially a factor analysis measuring how well two records match on all the criteria provided by the client, relative to the other records within the set of potential matches.

In crafting any matching algorithms, there is an important trade-off between precision and coverage. If the priority is to avoid false positive matches, then in cases where the correct match is ambiguous, one can be cautious by not matching the record at all. Catalist tells us that they favor precision over coverage in this way, as it is the priority of most of their campaign clients to avoid mistaking one voter for another. In contrast, matching algorithms that serve the military intelligence community (another high-volume user of matching programs) tend to favor coverage over precision, as the goal there is to identify potential threats.

One way to confirm the accuracy of Catalist's matching procedures is simply to highlight results from the MITRE name-matching challenge. According to Catalist, this is the one independent name-matching competition that allows companies to validate the quality of their matching

⁶The NES validation is problematic in another way too. It is based on a monthly panel study with attrition of over one-third of respondents from the initial interview to the November 2008 survey. Our validation was based on a standard pre-post-election panel with significantly less attrition.

⁷It is worth noting that we offered to match the 2008 NES with Catalist records, but the offer was declined. We understand that the NES is considering using commercial validation in the future.

procedure by participating in a third-party exercise. Participating teams are given two databases that mimic real-world complexities in having inconsistencies and alternate forms. One database might have typos, nicknames, name suffixes, etc., while the other does not. Of the forty companies competing, Catalist came in second place, above IBM and other larger companies.⁸ Their success in this competition speaks to their strength relative to the field.

Apart from the MITRE challenge, we asked Catalist to participate in our own validation exercise prior to matching the CCES to registration records. In 2009, along with our colleagues Alan Gerber and David Doherty, we conducted an audit of registration records in Florida and in Los Angeles County, California (Ansolabehere et al. 2010). We retrieved registration records from the state or county governments and sent a mail survey to a sample of registrants. After the project with our colleagues was complete, we sent to Catalist the random sample of voter records in Florida. We gave Catalist some but not all information from the voter's record. Once Catalist merged our reduced file with their database, they sent us back detailed records about each person. We then merged that file with our original records so that we can measure the degree of consistency. For example, there is a racial identifier in the voter file in Florida and Catalist sent us back a list of voters with race that they retrieved from their database. We did not give Catalist our record of the voter's race, so we know they could not have matched based on this variable. If Catalist identified the correct person in a match, then the two race fields should correspond. Indeed, in 99.9% of the cases, the two race fields match exactly ($N = 10,947$). For the exact date of registration, another field we did not transmit to Catalist, the two variables correspond to the day in 92.4% of the cases. For vote history, given that our records from the registration office identified an individual as having voted in the 2006 general election, 94% of Catalist's records showed the same; given that our record showed a person did not vote, 96% of Catalist's records showed the same. Of course, with vote history and registration date, the mismatches are very likely to be emanating from changes to records by the registration office rather than to false positive matches, as these two fields are more dynamic than one's racial identity.^{9,10}

The dynamic nature of registration records themselves points to a potential problem inasmuch as the registration data at the root of vote validation are imperfect. Registration list management is decentralized, voters frequently move, jurisdictions hold multiple elections every year—all these contribute to imperfect records. Of course, the records have improved greatly since the NES conducted its vote validations by visiting, in person, county election offices, and sorting through paper records. Not only has technology dramatically changed, but also the laws governing the maintenance of registration lists have changed. In 1993, Congress passed the National Voter Registration Act (NVRA), which included provisions detailing how voter records must be kept “accurate and current.” The Help America Vote Act (HAVA) of 2002 required every state (except North Dakota, which has no voter registration) to develop a “single, uniform, official, centralized, interactive computerized statewide voter registration list defined, maintained, and administered at the State level.”¹¹ The 2006 election was the first election by which states were required to have their databases up and running. It was not until after 2006 that Catalist and other data vendors were able to assemble national voter registration files that they could update regularly. The 2008 election is really the first opportunity to validate vote reports nationally using digital technologies.

Since the recent improvements in voter registration management, we have conducted several studies of registration list quality (see Ansolabehere and Hersh 2010; Ansolabehere et al. 2010). We have found the lists to be of quite high quality, of better quality in fact than the U.S. Census Bureau's mailing lists and other government lists. Two observations ought to be made about

⁸“Catalist Takes Second Place in MITRE Multi-Cultural Name Matching Challenge,” Catalist LLC, Press Release, October 6, 2011. For more information about the competition, consult “Conclusion of First MITRE Challenge Brings New Way to Fast-Track Ideas,” MITRE, Press Release, December 14, 2011.

⁹It is customary for registration offices to change a voter's date of registration if something about the registration record has been altered (e.g., a change in party affiliation or surname).

¹⁰Catalist's matching scheme is also validated in the sense that dozens of campaign clients rely on Catalist's matches to perform their voter contact programs. Catalist receives constant feedback about its models as the models are put to practical use on a daily basis. Catalist's success as a company is a market-based validation of their matching algorithm.

¹¹http://www.fec.gov/hava/law_ext.txt.

imperfections in voter registration records. First, the imperfections that have been found by us as well as by advocacy groups such as the Pew Center on the States are identified by comparing raw records to Catalist's clean records. Thus, working with a company like Catalist corrects for many issues with raw files such as purges, movers, and deceased voters. Second, in estimating the correlates of misreporting by comparing survey records to official records, errors on the registration files are likely to be random and simply contribute to measurement error. Consider the statistics cited above from our Florida validation: in the few cases in which Catalist's vote history did not correspond to the record we acquired from the state, the errors went in both directions (i.e., given our records showed a person voted, 6% of Catalist's records showed a nonvoter, and given our records showed a nonvoter, 4% of Catalist's records showed a voter). But when we turn to the data analysis and find that only validated nonvoters misreport their behavior and that misreporters follow a particular demographic profile, it is clear that there is far more to misreporting than measurement error.

As we turn to the data analysis and to the comparison of the NES vote validations of the 1980s with the 2008 CCES validation, we emphasize the differences between the surveys and advancements in validation methodology. We are comparing in-person cluster-sampled surveys validated by in-person examination of local hard-copy registration records with an online, probability-sampled, fifty-state survey validated by the electronic matching of survey identifiers to a commercial database of registration information collected from digitized state and county records. Though there are advantages to each of these survey modes (in-person versus online), it is fairly uncontroversial to assert that the quality of the data emanating from election offices and the technology of matching survey respondents to official records have improved dramatically in the past two decades. Here, we leverage these improvements to more accurately depict the correlates of misreporting and the true correlates of voting.

4 Data Analysis

4.1 *Reported, Misreported, and Validated Participation*

Before examining the kinds of respondents whose reported turnout history is at odds with their validated turnout history, we start with a basic analysis of aggregate reported vote and validated vote rates in the 2008 CCES and the 1980–1988 NES Presidential election surveys. As Cassel (2003) notes in reference to the NES validation studies, misreporting patterns are slightly different in Presidential and midterm election years, so it is appropriate to compare the 2008 survey with other Presidential year surveys. The statistics from the four Presidential election years can be found in Table 1. All statistics are calculated conditional on a nonmissing value for both the reported vote and validated vote. In other words, respondents who did not answer the question of whether they voted or for whom an attempt at validation was not possible are excluded.¹² It is important to note that in 1984 and 1988, the NES did not seek to validate respondents who claimed they were not registered to vote. Here, we treat these individuals as validated nonvoters, as is the convention. Given that these individuals admitted to being nonregistrants, we assume their probability of actually being registered or actually voting approaches zero.

The first two rows of data show the “true” turnout rates in each of the four elections. The first statistic is calculated by Michael McDonald (2011) as the number of ballots cast for President divided by the Voting Eligible Population. The second statistic is the Current Population Survey's estimate of turnout among the citizen population. Notice that CPS shows turnout rates of citizens to be similar across the election years under study, while McDonald's adjustments reveal that turnout of the eligible population was five to ten percentage points higher in 2008 than in the 1980s elections. While both the CCES and NES survey seek to interview U.S. adults, including noncitizens, it is probably the case that both survey modes face challenges in interviewing noncitizens and others who are not eligible for participation due to restrictions on felons and

¹²In the NES validations, no validation procedure was possible for respondents whose identifying information (e.g., their name) was not known or whose local election office would not grant access to the NES validators.

Table 1 Reported and validated vote rates in 1980, 1984, 1988, and 2008

Row		CCES 2008	NES		
			1988	1984	1980
	<i>“True” turnout rates</i>				
1	Among VEP (McDonald)	61.6	52.8	55.2	54.2
2	Among Citizen Pop. (CPS)	63.6	62.2	64.9	64.0
	<i>Survey turnout rates</i>				
3	Reported vote	84.3	69.5	73.7	71.5
4	Validated vote	68.5	60.1	63.8	61.2
5	Pr(Report vote valid vote)	99.1	98.8	99.8	99.4
6	Pr(Report vote valid not vote)	52.0	25.3	27.7	27.4
7	Pr(Valid vote report vote)	80.6	85.5	86.4	85.1
8	Observations	26,181	1718	1962	1279

Note. CPS statistics refer to Current Population Survey, November 2008, and Earlier Reports, “Table A-1. Reported Voting and Registration by Race, Hispanic Origin, Sex and Age Groups: November 1964 to 2008,” U.S. Census Bureau, July 2009 Release. VEP abbreviates the Voting Eligible Population.

ex-felons and the inaccessibility of the overseas voting population. As a result, McDonald’s eligible population figures are probably the superior baseline to gauge turnout rates in the surveys.

A comparison of rows 4 and 1 reveals the degree of sample selection bias in each survey, whereas a comparison of rows 4 and 3 reveals the degree of misreporting bias. As the difference between the validated vote rate and the VEP vote rate is smaller in the CCES than the NES, we see that sampling bias is greater in the in-person NES survey. Conversely, the differences between the reported and validated vote rates in the surveys reveal that misreporting bias is larger in the CCES. Row 5 shows that across surveys nearly every respondent who was validated as having voted reported that they voted. However, a large number of validated nonvoters also reported that they voted. In the NES, row 6 shows that just over a quarter of validated nonvoters claim they voted. In the CCES, half of nonvoters claim they voted.

What explains the increased level of misreporting in the CCES? One explanation might be that people are generally more comfortable lying on surveys now than they were a quarter-century ago; however, for lack of validation studies in the intervening years, this is not an hypothesis easily tested. Another explanation could be that 2008 presented a unique set of election-related circumstances that generated a spike in misreporting. This explanation is not plausible since we found a high rate of misreporting in the 2006 CCES (see Ansolabehere and Hersh 2011a) and we have preliminary results from our 2010 CCES-Catalist validation with a similar effect.¹³ Another explanation relates to Web surveys: If participants in Web surveys tend to be more politically knowledgeable than those participating in in-person surveys like the NES, and if politically knowledgeable people over-report their voting as the NES validation studies have shown, then we may have an explanation for the overreporting bias in the CCES as compared with the NES. To evaluate this claim, we will examine the correlates of misreporting below.

Before observing the correlates of misreporting, we take one other introductory cut at the data. We show in Fig. 1 the relationship between the reported vote rate and the validated vote rate for CCES respondents by state. In contrast to the NES validations that due to a cluster sampling scheme only validated vote reports in thirty to forty states, the CCES has been matched in fifty states plus Washington, DC. The only state not shown in Fig. 1 is Virginia, because Virginia does not permit out-of-state entities to examine its vote history data. Other pieces of the Virginia registration records, however, will be analyzed below. Immediately, the unusually low validated vote

¹³As a midterm election, misreporting is different than in Presidential years (Cassel 2003); our initial estimate from 2010 is that 35%–40% of validated nonvoters reported voting—still a considerably higher number than in the NES validation studies.

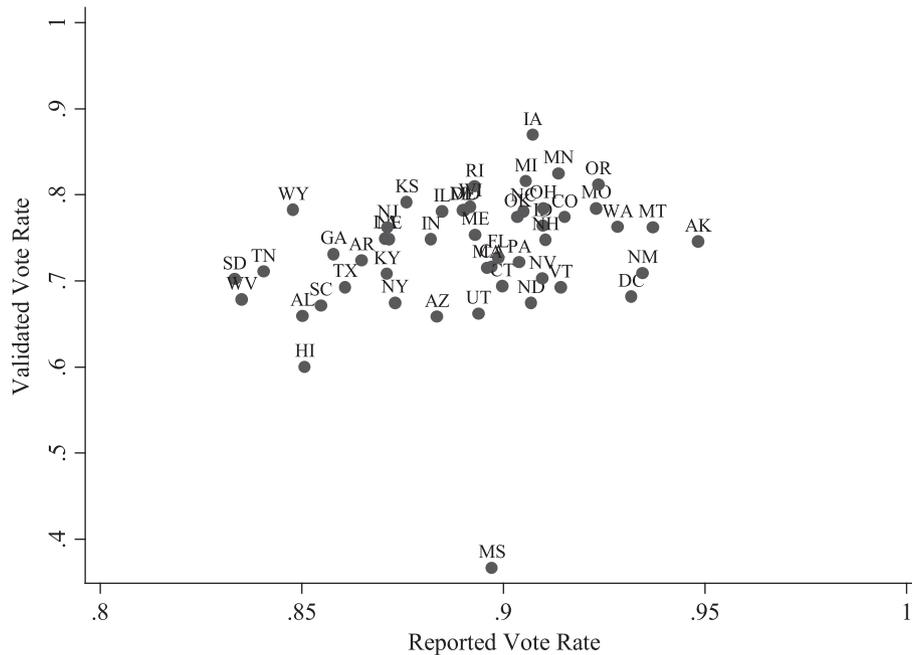


Fig. 1 State reported vote rates and validated vote rates. Reported voting and validated voting rates by state in the 2008 CCES matched to Catalist's national voter database.

rate in the state of Mississippi is the distinctive feature of Fig. 1. It turns out that Mississippi has, by far, the worst record in keeping track of voter history. This does not mean that the Mississippi voter file is generally of poor quality or that the matching of respondents to public or commercial records is bad in Mississippi, but simply that when it comes to correctly marking down which registrants voted in any given election, Mississippi is nearly twice as bad as any other state. Using a county-level measure of vote history discrepancies, we can condition our analyses on the quality of vote history data, as we will do below. Mississippi aside, there are no other major outliers in Fig. 1. Of course, states like Wyoming, Hawaii, and Alaska have smaller sample sizes and so there is variance due to state population size in the figure, but in general the validated vote rate ranges within a narrow band of twenty percentage points, just as the state-by-state official turnout rates do (McDonald 2011). And the reported vote rates in the states are in a narrow ten-percentage-point band.

4.2 Correlates of Vote Misreporting

Table 2 estimates the correlates of vote overreporting in the CCES and the 1980, 1984, and 1988 combined NES. The dependent variable is reported turnout, and only validated nonvoters are included. Because of the small NES sample sizes, the three NES surveys as combined in the NES cumulative data file are estimated together with year fixed-effects. To the best of our ability, we calibrate the independent variable measures to be as comparable as possible between the CCES and NES. Coding details and summary statistics are contained in the online appendix. The independent variables included in the model are a five-category education measure, a four-category income measure, and two indicator variables for race—African American and other non-White. In the combined NES sample, there are only forty-seven non-Black minorities who said they voted but were validated as nonvoting, and thus we do not cut up the racial identifier more fine-grained than that. Other independent variables estimated are indicator variables for married voters, women, recent movers (residing two years or fewer in current home), and four indicator variables representing age groups. A measure of partisan strength ranges from pure independents (0) to strong partisans (3). Finally, we include measures for the frequency of

Table 2 Regression models of overreporting compared across surveys

<i>Dep Var:</i>	<i>CCES</i>	<i>NES</i>
<i>Reported vote</i>	<i>2008</i>	<i>1980–1984–1988</i>
<i>Indep. Vars.:</i>	$\hat{\beta}$	$\hat{\beta}$
Education	0.069** (0.058 to 0.079)	0.086** (0.064 to 0.109)
Income	0.056** (0.045 to 0.067)	0.052** (0.030 to 0.074)
Black	0.022 (−0.021 to 0.065)	0.065* (0.008 to 0.123)
Other non-White	−0.040** (−0.068 to −0.011)	−0.010 (−0.072 to 0.052)
Married	−0.006 (−0.027 to 0.015)	−0.049* (−0.093 to −0.005)
Church attendance	0.032** (0.022 to 0.041)	0.049** (0.029 to 0.068)
Age (years)		
25–34	−0.060** (−0.096 to −0.024)	0.023 (−0.036 to 0.082)
35–44	−0.024 (−0.061 to 0.014)	0.084* (0.018 to 0.149)
45–54	−0.075** (−0.112 to −0.039)	0.090* (0.011 to 0.168)
≥55	0.028 (−0.009 to 0.064)	0.128** (0.063 to 0.193)
Ideological strength	0.009 (−0.005 to 0.023)	0.004 (−0.033 to 0.041)
Female	−0.145** (−0.166 to −0.125)	−0.006 (−0.048 to 0.036)
Poly. interest	0.155** (0.144 to 0.166)	0.068** (0.047 to 0.089)
Partisan strength	0.066** (0.057 to 0.075)	0.055** (0.034 to 0.075)
Recent mover	−0.027* (−0.049 to −0.006)	−0.091** (−0.152 to −0.030)
Year		
1984		−0.012 (−0.063 to 0.039)
1988		−0.018 (−0.069 to 0.032)
Constant	−0.069** (−0.117 to −0.020)	−0.228** (−0.317 to −0.140)
Observations	6380	1633
R^2	0.357	0.179

Note. 95% confidence intervals are in parentheses. OLS regressions are shown. Logistic regression tables are available in an online appendix. Models of overreporting estimate reported voting among respondents who have been validated as nonvoters. Coding of variables for the NES and CCES can be found in the online appendix. ** $p < .01$, * $p < .05$.

reported church attendance, the degree of ideological extremism (i.e., a measure ranging from moderate to either very conservative or very liberal), and one's level of interest in public affairs.

Table 2 shows ordinary least squares regression coefficients. As the dependent variable (reported voting) is binary, we also include an online appendix where all tables are replicated using logistic regression and reporting standard errors rather than confidence intervals. A comparison of the coefficients of the CCES and NES models reveals a number of similarities between the estimates from two very different kinds of surveys validated in two different ways separated in time by over

two decades. Well-educated, high-income partisans who are engaged in public affairs, attend church regularly, and have lived in the community for a while are the kinds of people who misreport their vote experience in both cases. Consistent with the findings of Silver, Anderson, and Abramson (1986), Belli, Traugott, and Beckmann (2001), Bernstein, Chaha, and Montjoy (2001), and many others, it is the nonvoters who have the politically relevant resources to be engaged with politics who are most likely to report that they voted.

There are, however, some noteworthy differences between the two models. First, in the NES surveys it appears that misreporting increases with age. Interestingly, in 2008, the age groups most likely to misreport are the oldest cohort (ages 55 and up) and also the *youngest* cohort (which is the excluded age category in the model). To the extent that young people were particularly mobilized in the 2008 election, they might have been under unusually high pressure to be considered a voter and so it is understandable that their rate of misreporting is high in this election. In the NES regression in Table 2, misreporting was not a function of gender, but in the 2008 CCES it appears that men misreported more than women. Likewise, the relationship between race and misreporting is noticeably different in the NES and CCES, as is the magnitude of the effect on political interest. In spite of some differences, the upshot of Table 2 is the remarkable degree of consistency between the in-person NES validations of the 1980s and the digital CCES validation of 2008.

4.3 Accounting for the Quality of Registration Lists

In spite of the consistent patterns of misreporting across election years, there is yet a concern that validating voting records is fraught with error because official registration data are often messy and are constantly changing, and thus the matching procedure is unreliable. A specific concern is that states differ in the kinds of data they collect from registrants and in the quality of the record-keeping system, and thus there may be serious differences in the quality of matching across states. In this section, we attempt to parse out misreporting from registration quality issues that may generate inconsistencies between respondent reports and official records. In Table 3, we begin to reestimate the CCES coefficients in Table 2 in several ways.

The first column in Table 3 is the same as the CCES model in Table 2 and is repeated for the sake of easy reference. In the second model, we incorporate state fixed effects. If we test the equivalence of these two models with a likelihood ratio test, they are estimated to be distinct (chi-squared value: 118.7), indicating average state differences in the rate of misreporting, but notice that the coefficient estimates between the columns are indistinguishable. Thus, whatever difference in the quality or status of voting records across states, accounting for state-by-state differences does not alter the relationship between the individual-level variables and misreporting.

In the third column, we restrict the fixed-effects regression to only respondents whom Catalist identified. We, as well as Catalist, take individuals not found in the national database as unregistered. Suppose that many of these individuals are in fact registered but poor records or poor matching caused them to appear as unregistered. In column 3, we explore whether restricting the sample to only those found by Catalist (including registrants, former registrants, and unregistered voters who appear in consumer databases) alters the relationship between personal traits and misreporting. Obviously, since we are removing from the sample a large group of individuals thought to be unregistered, the relationship among some variables will change; but notice that on most of the key demographics, a similar pattern emerges as in the other estimations. Finally, the fourth column adds one individual-level variable to the model estimated in column 3—Catalist's score indicating how confident it is in the match between the survey and the voter database. Recall that Catalist used name, address, birth year, and gender to match respondents to voter files. For some respondents, there may be the possibility that two records on the voter file representing different people looked somewhat similar. The confidence score measures that uncertainty. A high score indicates higher confidence. Note that the coefficient on the score is not significant, and a likelihood ratio test comparing model 3 to model 4 shows that the two are not distinct.

Next we turn to three county-level indicators of registration list quality. The first estimates the percentage of records in a county that are considered by Catalist to be likely or probably

Table 3 Regression models incorporating state fixed-effects and individual-level match confidence control

<i>Dep Var:</i> <i>Reported vote</i> <i>Indep. Vars.:</i>	<i>Basic model</i> $\hat{\beta}$	<i>Add state</i> <i>fixed-effects</i> $\hat{\beta}$	<i>Restricted to</i> <i>matched Rs</i> $\hat{\beta}$	<i>Add indiv.-level</i> <i>conf. measure</i> $\hat{\beta}$
Education	0.069** (0.058 to 0.079)	0.068** (0.057 to 0.078)	0.083** (0.068 to 0.097)	0.082** (0.068 to 0.097)
Income	0.056** (0.045 to 0.067)	0.056** (0.045 to 0.067)	0.053** (0.038 to 0.068)	0.053** (0.038 to 0.067)
Black	0.022 (-0.021 to 0.065)	0.030 (-0.014 to 0.073)	0.054 (-0.006 to 0.113)	0.053 (-0.006 to 0.112)
Other non-White	-0.040** (-0.068 to -0.011)	-0.037* (-0.066 to -0.008)	-0.041 (-0.084 to 0.002)	-0.041 (-0.084 to 0.002)
Married	-0.006 (-0.027 to 0.015)	-0.006 (-0.027 to 0.016)	-0.002 (-0.031 to 0.026)	-0.002 (-0.031 to 0.026)
Church attendance	0.032** (0.022 to 0.041)	0.032** (0.023 to 0.041)	0.039** (0.026 to 0.051)	0.039** (0.026 to 0.051)
Age (years)				
25–34	-0.060** (-0.096 to -0.024)	-0.058** (-0.094 to -0.022)	-0.014 (-0.067 to 0.039)	-0.014 (-0.067 to 0.038)
35–44	-0.024 (-0.061 to 0.014)	-0.027 (-0.064 to 0.010)	-0.007 (-0.060 to 0.046)	-0.008 (-0.061 to 0.045)
45–54	-0.075** (-0.112 to -0.039)	-0.075** (-0.111 to -0.038)	-0.056* (-0.108 to -0.003)	-0.057* (-0.110 to -0.005)
≥55	0.028 (-0.009 to 0.064)	0.023 (-0.013 to 0.060)	0.025 (-0.027 to 0.077)	0.024 (-0.029 to 0.076)
Ideological strength	0.009 (-0.005 to 0.023)	0.009 (-0.005 to 0.023)	-0.001 (-0.020 to 0.018)	-0.001 (-0.020 to 0.018)
Female	-0.145** (-0.166 to -0.125)	-0.146** (-0.166 to -0.126)	-0.126** (-0.153 to -0.099)	-0.125** (-0.152 to -0.098)
Poly. interest	0.155** (0.144 to 0.166)	0.152** (0.141 to 0.164)	0.146** (0.131 to 0.162)	0.146** (0.131 to 0.161)
Partisan strength	0.066** (0.057 to 0.075)	0.065** (0.056 to 0.074)	0.070** (0.057 to 0.082)	0.070** (0.057 to 0.082)
Recent mover	-0.027* (-0.049 to -0.006)	-0.026* (-0.048 to -0.005)	0.017 (-0.012 to 0.046)	0.016 (-0.013 to 0.045)
Confidence				-0.254 (-0.625 to 0.116)
Constant	-0.069** (-0.117 to -0.020)	0.107 (-0.111 to 0.324)	-0.027 (-0.423 to 0.370)	0.187 (-0.317 to 0.692)
State fixed-effects?	No	Yes	Yes	Yes
Observations	6380	6380	3710	3710
R ²	0.357	0.369	0.362	0.362

Note. 95% confidence intervals are in parentheses. OLS regressions are shown. Logistic regression tables are available in an online appendix. ** $p < .01$, * $p < .05$.

deadwood. This captures records of people who are unrealistically old, people identified as deceased in the Social Security Death Index, or people who have moved or have not voted at their listed address in several years. The second county-level measure is based on the U.S. Post Office's Coding Accuracy Support System (CASS) that estimates the proportion of mailing addresses on the voter file thought to be undeliverable as addressed. A high rate of undeliverable addresses provides another signal that the election administration in the county may be inadequately managing records. The final measure is the absolute deviation between the number of registrants marked as having voted in the 2008 election and the number of votes officially counted in the 2008 election, divided by the official count in the county. Large deviations suggest a problem in the calibration between registration-listed vote history and actual voting patterns.

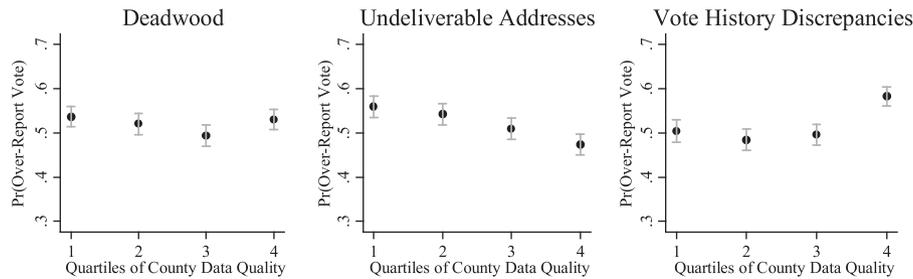


Fig. 2 Rates of misreporting by quality of registration lists in counties. Counties are divided into quartiles based on the quality of their records. Higher quartiles indicate counties with higher rates of deadwood, undeliverable address, or vote history discrepancies. 95% confidence intervals are shown.

In Fig. 2, respondents are binned into quartiles based on the quality of registration records in the counties in which they are registered. Higher quartiles indicate counties that performed worse on these measures. As the percentage of records considered to be deadwood increases in a county, the rate of respondents reporting voting when their records show they did not vote is unaffected. In counties with a high incidence of invalid registration addresses, misreporting is actually slightly but perceptibly lower than in better-quality counties. Finally, in the third measure—the one that measures vote history quality directly—there is a higher rate of apparent-misreporters in the counties in which official turnout was very different from turnout as calculated from voter files. In a few outlier counties, all or almost all registrants were marked by the election office as not having voted. Thus, it is not surprising that in these counties, there is a higher rate of what appears like misreporting. To put this in perspective, however, 75% of respondents live in counties where fewer than 2% of records indicate a discrepancy on vote history, and 97% of respondents live in counties where fewer than 10% of records have such a discrepancy. It is really only a few counties in which vote history discrepancies likely contribute to the appearance of misreporting.

There are a number of ways to treat these county-level measures, depending on the research purpose. For some purposes, it may be worth omitting survey respondents who live in the few counties with unreliable vote history data. For our purposes, we would like to study the extent to which individual-level correlates of misreporting differ in counties with more and less clean records. Suppose that respondents whose survey response do not match their public record disproportionately live in counties where registration records are unreliable. Also suppose that these places have a higher rate of citizens who are well-educated, high-income partisans who are engaged in public affairs, attend church regularly, and have lived in the community for a while. This is the sort of the scenario that could lead one to mistakenly infer that respondent characteristics rather than bad-quality records are at the heart of misreporting. On its face, this scenario seems unlikely to be true if we believe that the high-SES, politically active individuals are the same sorts of people who demand high-performing governments. But again, we must return to the recent NES validation project, which concluded that “difficulties in locating government records caused problems for attempts to measure and explain turnout” (Berent, Krosnick, and Lupia 2011, 62). These scholars found that when they relaxed their matching criteria and attempted to validate records that were less complete, and when they then predicted turnout, the correlates of participation differed from samples in which they only validated respondents with the cleanest records.

For Table 4, we first create a summary county-level statistic that is each county’s average value on the three measures of list quality described above. We then divide respondents into four quartiles based on their scores on this county measure, and we estimate OLS equations of misreporting on individual-level covariates in each county environment. Across all four groups, the coefficients are comparable. There is some fluctuation on the two race variables, though within the bounds of normal measurement error that can be expected when quartering a sample. A chow test statistic of 2.87 ($F(16, 6,335)$) indicates that coefficients in the different subsamples are statistically distinct from the grouped sample. Nevertheless, the close resemblance of coefficients across county types is

Table 4 Comparison of misreporting in counties grouped by the quality of registration records

<i>Dep Var:</i> <i>Reported vote</i> <i>Indep. Vars.:</i>	<i>All counties</i> $\hat{\beta}$	<i>Best quality</i> <i>counties</i> $\hat{\beta}$	<i>Next best</i> <i>counties</i> $\hat{\beta}$	<i>Second worst</i> <i>counties</i> $\hat{\beta}$	<i>Worst counties</i> $\hat{\beta}$
Education	0.069** (0.058 to 0.079)	0.067** (0.046 to 0.089)	0.080** (0.058 to 0.102)	0.066** (0.044 to 0.088)	0.064** (0.045 to 0.082)
Income	0.056** (0.045 to 0.067)	0.047** (0.025 to 0.070)	0.046** (0.023 to 0.068)	0.069** (0.046 to 0.092)	0.067** (0.047 to 0.086)
Black	0.022 (-0.021 to 0.066)	-0.042 (-0.141 to 0.057)	-0.017 (-0.107 to 0.073)	0.044 (-0.037 to 0.125)	0.088* (0.008 to 0.168)
Other non-White	-0.040** (-0.068 to -0.012)	-0.013 (-0.064 to 0.039)	-0.080** (-0.139 to -0.020)	-0.080* (-0.145 to -0.016)	0.000 (-0.057 to 0.057)
Married	-0.006 (-0.027 to 0.015)	-0.035 (-0.079 to 0.009)	0.019 (-0.025 to 0.064)	0.028 (-0.016 to 0.072)	-0.034 (-0.072 to 0.004)
Church attendance	0.031** (0.022 to 0.040)	0.036** (0.017 to 0.055)	0.019* (0.000 to 0.038)	0.037** (0.018 to 0.057)	0.032** (0.016 to 0.049)
Age (years)					
25-34	-0.059** (-0.095 to -0.023)	-0.053 (-0.126 to 0.020)	-0.028 (-0.101 to 0.045)	-0.088* (-0.163 to -0.013)	-0.068 (-0.138 to 0.001)
35-44	-0.023 (-0.060 to 0.014)	0.000 (-0.075 to 0.076)	-0.028 (-0.105 to 0.050)	-0.053 (-0.128 to 0.023)	-0.022 (-0.093 to 0.048)
45-54	-0.075** (-0.111 to -0.038)	-0.093* (-0.168 to -0.017)	-0.090* (-0.165 to -0.015)	-0.106** (-0.182 to -0.029)	-0.035 (-0.104 to 0.035)
≥55	0.029 (-0.008 to 0.066)	0.009 (-0.066 to 0.085)	0.011 (-0.065 to 0.087)	0.022 (-0.055 to 0.098)	0.055 (-0.013 to 0.123)
Ideological strength	0.009 (-0.005 to 0.023)	0.008 (-0.022 to 0.037)	-0.021 (-0.051 to 0.008)	0.013 (-0.017 to 0.042)	0.032* (0.007 to 0.056)
Female	-0.145** (-0.166 to -0.125)	-0.161** (-0.203 to -0.119)	-0.123** (-0.165 to -0.081)	-0.149** (-0.192 to -0.106)	-0.147** (-0.184 to -0.111)
Poly. interest	0.155** (0.144 to 0.167)	0.160** (0.136 to 0.184)	0.164** (0.140 to 0.188)	0.150** (0.126 to 0.174)	0.148** (0.127 to 0.169)
Partisan strength	0.066** (0.057 to 0.075)	0.068** (0.048 to 0.087)	0.070** (0.051 to 0.088)	0.060** (0.042 to 0.079)	0.063** (0.046 to 0.079)
Recent mover	-0.028* (-0.049 to -0.006)	-0.027 (-0.072 to 0.018)	-0.033 (-0.078 to 0.012)	-0.006 (-0.050 to 0.039)	-0.042* (-0.081 to -0.003)
Constant	-0.069** (-0.117 to -0.020)	-0.052 (-0.155 to 0.051)	-0.062 (-0.162 to 0.037)	-0.103* (-0.204 to -0.002)	-0.067 (-0.156 to 0.022)
Observations	6367	1519	1510	1439	1899
R ²	0.357	0.341	0.345	0.385	0.373

Note. 95% confidence intervals are in parentheses. OLS regressions are shown. Logistic regression tables are available in an online appendix. ** $p < .01$, * $p < .05$.

clear. Of particular importance is the second model, which is restricted to only the counties with the highest-rated registration records. In these counties, there is little reason to believe that misreporting could be a function of poor records, yet here (like everywhere) the same type of citizen tends to over-report voting.

While a small portion of misreporting may be attributable to matching survey respondents to records that are sometimes erroneous, imperfections in the validation process are not the main story of misreporting. In fact, compared with the demographic variables, the detailed state-level, county-level, and individual-level measures of validation quality explain little variance in estimating a model of misreporting (compare, for instance, the R^2 across models in Tables 3 and 4). This represents the first attempt to take seriously concerns about validation methodology by controlling both for the quality of the matching procedure at the level of individual respondents and by employing county-level measures of election administration quality. None of these additions alter the results articulated by Silver, Anderson, and Abramson (1986) and others that over-reporters are individuals who look much like voters in the demographics and attitudes but who did not actually vote.

4.4 *Validating Other Attributes and Behaviors*

The Catalist validation project enables us to validate more than just voting, and here we validate four other self-reported survey items: race, party affiliation, method of voting, and registration. Examining misreporting on these other items allows us to test whether misreporting is likely to be related to social desirability. As mentioned in Section 3, this is also yet another way to validate the quality of Catalist's vote validation procedure. When reporting about one's party, race, or vote method, we have little reason to believe that respondents will feel compelled to lie. If the same types of people misreport on these items as on voting, then something other than social desirability may be at the root of misreporting. It might imply that people are clicking through responses without paying much attention, that the matching algorithm was not successful, or that something other than lying (like memory) is at the root of misreporting. If, however, respondents tend to misreport on items like voting but not on mundane items, then the social desirability hypothesis seems more compelling.

In order for Catalist to match CCES respondents to voter file records, Polimetrix transferred to Catalist the names, addresses, birth years, and genders of respondents. This was all the information that Catalist used to identify voters. Table 5 shows the rate of misreporting on other items we could verify with official records: party, race, vote method, and registration. The CCES asked about respondents' party registration in addition to party affiliation. Official records of party registration exist in about half of the states. We matched reported racial identity to the race listed on the voter files in the southern states that ask registrants about their racial identities. For both party and race, validated Democrats, Republicans, Blacks, and Whites reported the same response as listed in the record at a rate of 93%–95%. There is less consistency with the "other" categories, which might be explained either by the way states and counties keep track of smaller population groups or by respondents with more ambiguous racial or partisan identities having some inconsistencies in their reported traits.

In the next part of Table 5, we observe the relationship between the method of voting reported by a voter and the method of voting recorded by an election office. In many states, voters have the option of voting ahead of Election Day, either by submitting a mail ballot or by voting early in person. Ninety-six percent of voters who were validated as voting early or by mail reported voting this way. A somewhat lower 85% of those validated as voting at the polls reported voting at the polls. It is not immediately obvious why validated early/absentee voters are more consistent with the official records than precinct voters, but we will return to this question in a moment.¹⁴

¹⁴Validated nonvoters who report a vote method are distributed across methods of voting similarly to validated voters. In other words, about 60% of validated voters report that they voted at the polls and about 60% of validated nonvoters report that they voted at the polls. This similarity supports the notion that misreporters are people who generally vote or know about voting and therefore provides an answer that matches their actually voting peers.

Table 5 Misreporting on party registration, racial identity, vote method, and registration status

	<i>Percent</i>
Party	
Pr(Rep. D Val. D)	94.6
Pr(Rep. R Val. R)	93.3
Pr(Rep. Other Val. Other)	73.9
Pr(Rep. D Val. R)	2.9
Pr(Rep. R Val. D)	2.0
Pr(Rep. Other Val. D, R)	3.6
Pr(Rep. D, R Val. Other)	26.1
Observations	11,292
Race	
Pr(Rep. B Val. B)	95.1
Pr(Rep. W Val. W)	94.3
Pr(Rep. Other Val. Other)	92.7
Pr(Rep. B Val. W)	0.7
Pr(Rep. W Val. B)	1.7
Pr(Rep. Other Val. B, W)	9.2
Pr(Rep. B, W Val. Other)	7.3
Observations	4540
Vote method	
Pr(Rep. Polls Val. Polls)	85.1
Pr(Rep. Early/Abs. Val. Early/Abs)	96.1
Pr(Rep. Polls Val. Early/Abs)	3.9
Pr(Rep. Early/Abs Val. Polls)	15.0
Observations	19,145
Registration	
Pr(Rep. Reg Val. Reg)	97.8
Pr(Rep. Not Reg. Val. Not Reg.)	35.8
Pr(Rep. Reg. Val. Not Reg.)	64.2
Pr(Rep. Not. Reg Val. Reg.)	2.3
Observations	26,864

Note. Rep. abbreviates reported; Val. abbreviates validated. Party registrants are separated into Democrats (D), Republicans (R), and independents and third-party registrants (Other). Racial groups are separated into Whites (W), Blacks (B), and others. Voters are separated into those who cast a ballot on Election Day at the polls (polls) and those who cast a ballot by mail or early in person (Early/Abs.).

At the bottom of Table 5, we validate reported registration status. We include as validated registrants active as well as inactive registered voters, plus those who were registered at an old address but not at their current one. We see a pattern that is very similar to the validated vote statistics in Table 1. Nearly all (98%) respondents validated as registered reported being registered, whereas a substantial percentage of validated nonregistrants (64%) reported that they were registered. Like with voting, there seems to be a trend of nonparticipants claiming to be participants.

To get a better handle on how misreporting varies with the subject of the survey question, we show in Table 6 four models of misreporting. First, we repeat our standard model of vote misreporting from Tables 2 and 3. We also show parallel models for registration, party, and vote method.¹⁵ For registration, we measure who among validated nonregistrants reported that they were registered. For party, we measure who among validated third-party or independent voters

¹⁵We do not estimate a model of race misreporting because very few respondents reported a different race than their registration record showed and because the sample size is already restricted on account of race data only being available in a few states.

Table 6 Regression models of overreporting about voting, registration, vote method, and party affiliation

<i>Dep Vars:</i> <i>Indep. Vars.:</i>	<i>Vote misreporting</i> $\hat{\beta}$	<i>Registration misreporting</i> $\hat{\beta}$	<i>Absentee/early misreporting</i> $\hat{\beta}$	<i>Party affiliation misreporting</i> $\hat{\beta}$
Education	0.069** (0.058 to 0.079)	0.057** (0.046 to 0.069)	0.017** (0.010 to 0.023)	-0.026** (-0.043 to -0.008)
Income	0.056** (0.045 to 0.067)	0.033** (0.021 to 0.046)	-0.005 (-0.013 to 0.003)	-0.010 (-0.031 to 0.011)
Black	0.022 (-0.021 to 0.065)	-0.028 (-0.077 to 0.021)	0.058** (0.030 to 0.086)	0.092* (0.001 to 0.183)
Other non-White	-0.040** (-0.068 to -0.011)	-0.045** (-0.076 to -0.014)	-0.022 (-0.044 to 0.001)	0.008 (-0.040 to 0.057)
Married	-0.006 (-0.027 to 0.015)	-0.005 (-0.029 to 0.019)	-0.008 (-0.023 to 0.007)	0.009 (-0.029 to 0.048)
Church attendance	0.032** (0.022 to 0.041)	0.028** (0.018 to 0.039)	0.004 (-0.001 to 0.010)	0.022** (0.006 to 0.037)
Age (years)				
25–34	-0.060** (-0.096 to -0.024)	-0.026 (-0.067 to 0.015)	-0.057** (-0.088 to -0.025)	0.094** (0.024 to 0.164)
35–44	-0.024 (-0.061 to 0.014)	-0.028 (-0.070 to 0.014)	-0.040* (-0.071 to -0.009)	0.118** (0.047 to 0.190)
45–54	-0.075** (-0.112 to -0.039)	-0.054* (-0.095 to -0.013)	-0.017 (-0.047 to 0.014)	0.122** (0.051 to 0.193)
≥55	0.028 (-0.009 to 0.064)	0.001 (-0.041 to 0.042)	0.046** (0.016 to 0.076)	0.098** (0.028 to 0.168)
Ideological strength	0.009 (-0.005 to 0.023)	0.002 (-0.014 to 0.018)	-0.004 (-0.012 to 0.005)	-0.028* (-0.052 to -0.004)
Female	-0.145** (-0.166 to -0.125)	-0.137** (-0.160 to -0.114)	0.002 (-0.011 to 0.015)	-0.024 (-0.060 to 0.011)
Poly. interest	0.155** (0.144 to 0.166)	0.132** (0.119 to 0.145)	0.012* (0.002 to 0.022)	-0.032** (-0.056 to -0.009)
Partisan strength	0.066** (0.057 to 0.075)	0.055** (0.045 to 0.065)	0.001 (-0.005 to 0.008)	0.233** (0.215 to 0.250)
Recent mover	-0.027* (-0.049 to -0.006)	-0.010 (-0.034 to 0.015)	0.030** (0.014 to 0.046)	0.019 (-0.022 to 0.059)
Constant	-0.069** (-0.117 to -0.020)	0.198** (0.143 to 0.253)	0.093** (0.050 to 0.136)	0.008 (-0.092 to 0.108)
Observations	6380	4552	12,515	1797
R^2	0.357	0.297	0.017	0.299

Note. 95% confidence intervals are in parentheses. OLS regressions are shown. Logistic regression tables are available in an online appendix. ** $p < .01$, * $p < .05$.

reported they were registered as a Democrat or Republican. For vote method, we measure who among validated precinct voters reported that they voted absentee or early. For party and vote method, these are the groups for which there seems to be more misreporting than is typical of other groups on these items. Based on the data in Table 5, there seems to be a general level of misreporting at 5%–10% that might be attributable to bad record-keeping or incorrect responses. The misreporting captured in the models in Table 6 identifies the survey questions that attract more misreporting than that.

The results in Table 6 demonstrate that the same characteristics that predict misreporting about voting also predict misreporting about registration, but do not predict other forms of misreporting in the same way. For instance, it is not the case that education, income, church attendance, gender, political interest, or ideology are anywhere as predictive of misreporting about party and vote method as they are about voting and registration. Misreporters of party and vote method have

interesting characteristics in their own right, as indicative by the coefficients in Table 6, and we could tell stories about why certain types of voters tended to misreport these behaviors in 2008. But Table 6 indicates that the standard correlates of vote and registration misreporting are not at work in the same way with respect to behaviors or characteristics that do not have a socially desirable response.

The upshot of this validation of survey reports about party, race, and voting method is that misreporting on voting and registering is a different phenomenon than typical inconsistencies between surveys and official public records. When a survey participant is validated as White or Black, Democratic or Republican, in about 94% of cases the respondent will report just as the record shows. However, this is not the case when respondents are asked whether they are registered to vote or whether they voted in a recent election. On these items, a very predictable set of participants—those who are of high socioeconomic status, engaged in their communities, and interested in politics—routinely misrepresent their officially recorded behavior. That the same set of voters misreports in every validated political survey, that by comparison measures of election administration quality do not predict nearly as much misreporting as demographic and attitudinal descriptors, and that survey reports about items unrelated to voting exhibit a high level of consistency with matched official records advances the theory that misreporting is not about the quality of survey matching but about a certain set of nonvoting individuals who like to think of themselves as voters.

5 So Who *Really* Votes?

The most important consequence of misreporting is that since misreporters look like voters, comparing reported voters and nonvoters based on surveys exaggerates the differences between the two groups. This section reveals the extent of this problem. Similar to voters, misreporters are disproportionately well-educated, wealthy, partisan, and interested in politics. Thus, when survey researchers compare voters and nonvoters based on reported turnout, they count as voters a particularly engaged set of individuals who actually did not vote. If these nonvoting, but engaged, individuals are outed through validation, nonvoters and voters begin to look much more similar to one another than they would just by studying reported behaviors.

Reported nonvoters are a distinct set of people who not only fail to vote but also feel little social or psychological imperative to lie about not voting (or maybe they just find lying to be unusually distasteful). Validated nonvoters include all of these reported nonvoters plus an additional set of people who are very much engaged with politics but who feel some need to misrepresent their record. This latter group is, in fact, larger than the former group in the 2008 CCES. When combined together, admitted nonvoters and misreporting nonvoters look less different in their demographics and attitudes from actual voters.

To briefly examine the correlates of reported turnout and the correlates of validated turnout, in Fig. 3 we study thirteen demographic traits and calculate the percentage of respondents possessing those traits given that they are reported voters, reported nonvoters, valid voters, and valid nonvoters. Each of these variables are coded as indicator variables for easy interpretation. We subtract the mean value for nonvoters from the mean value for voters, build a confidence interval around the difference (though the intervals are too narrow to appear in the figure), and display the differences in Fig. 3. For example, reported voters are twenty-two percentage points more likely to have a bachelor's degree than reported nonvoters. But valid voters are only ten percentage points more likely to have a degree than valid nonvoters. It is clear that using reported vote data exaggerates the extent to which voters appear to be different from nonvoters especially with respect to education, income, political interest, and partisanship. Gender is another interesting variable here. Studying reported voters and nonvoters, it looks as if men voted more than women; but in the validated vote model, the reverse is apparent, and we know it to be correct that women voted at higher rates than men in 2008 (see Ansolabehere and Hersh 2011b, a corrective to earlier work on gender and participation, such as Burns, Schlozman, and Verba 2001). If we estimate a multivariate regression model with these sorts of demographics using validated vote as a dependent variable rather than reported vote, the R^2 value is cut nearly in half, indicating that these personal traits that

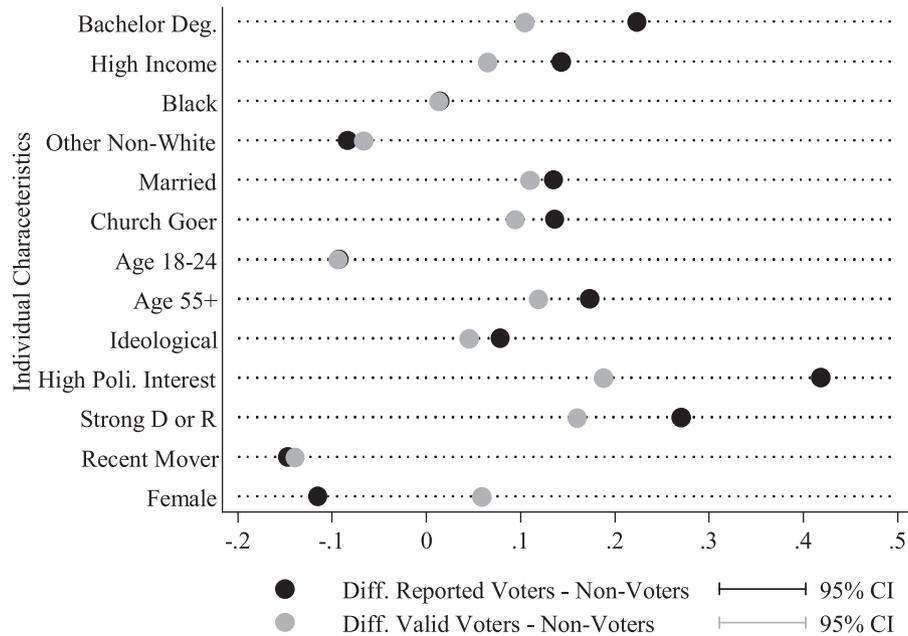


Fig. 3 Demographic and attitudinal differences between reported voters and nonvoters and between validated voters and nonvoters. 95% confidence levels are displayed, but in all cases are smaller than the width of the dots and so cannot be seen. Thus, dots that are separated indicate statistically significant differences. All variables are converted to indicator measures. *High Income* is an indicator for those reporting a family income \geq \$100K; *Church Goer* is an indicator for those reporting attending services at least weekly; *Ideological* and *Strong D or R* are indicators for those reporting to be very conservative/liberal and strong Republicans/Democrats, respectively.

dominate resource-based models of participation explain much less about voting than one would gather from looking at reported vote models alone.

The result that voters are less different from nonvoters than would be observed from survey responses is important because it speaks to the issues of equality and representation articulated in classic works by Verba, Schlozman, and Brady (1995), Rosenstone and Hansen (1993), and others. These scholars follow in a line of democratic theorists and social scientists who are concerned about how well citizens are represented by the cohort of engaged participants. In a society where political participation is voluntary, it is possible that those who volunteer to take an active role have different preferences for government as compared to those who opt out of participation. If the activists vote in such a way that serves their particular interests and ignores the interests of nonvoters, this might be concerning. Furthermore, if participation is costly such that citizens with fewer resources are unable to take part and their interests go unrepresented, this too may be cause for concern.

The evidence in Fig. 3 tempers the concerns of past scholarship built solely on the reported survey data regarding the characteristics of voters and nonvoters and the issue of representation. Looking at survey data alone, as Verba, Schlozman, and Brady (1995), Rosenstone and Hansen (1993), and many others do, voters and nonvoters are very different in their demographic and attitudinal traits than validated records reveal. In reality, voters and nonvoters are different from one another, but not nearly as different as is generally estimated by survey data. In every election cycle, a great number of individuals who care about politics enough to vote and have the resources enough to vote do not end up voting. For lack of time or effort, they join millions of Americans in abstaining from election participation. But these individuals (who make up about 15% of the public, according to the 2008 CCES) report in surveys that they did vote. All the evidence points to the conclusion that these individuals are not voters, and for political scientists to truly

grapple with the equality and representation of voters and nonvoters, misreporters must be identified and treated as nonvoters.

6 Conclusion

The overestimation of turnout in public opinion surveys is due in part to sample selection bias and in part to vote misreporting. In the NES of the 1980s, sample selection was a bigger contributing factor than misreporting; in the 2008 CCES, misreporting was a bigger contributing factor than sample selection. In both cases, a particular set of respondents have been found to consistently misreport. While a small amount of misreporting can be explained by measures of registration list quality and a small amount may be due to random error or mismatched records, by orders of magnitude these factors explain far less about misreporting than the simple demographics that identify the well-educated, high-income, partisan, politically active, church-attending respondents who lie about their participatory history.

The dramatic effect of misreporting on models of participation demands a renewed effort at theory-building. Sociodemographic and political resources do not explain all that much about why certain people vote and others do not. These variables, as shown in Fig. 3, simply perform the dubious function of identifying survey respondents who think of themselves as voters. The resources model explains less than half of the difference between voters and nonvoters as is conventionally thought. For those who worry about the causes and consequences of actual political participation, we need models that take us further along.

This research should thus spur more survey validation in the future and also spur attempts to improve methodological issues of biased samples and misreporting. Though the main focus in this article has been on misreporting, sample selection is an equally important problem and also demands attention. The evidence of this analysis only comes from one survey, so additional validations of other surveys will be necessary to draw firmer conclusions about the persistent effects of misreporting. Through partnering with a commercial vendor and leveraging new data and technology, we have found survey matching to be relatively inexpensive and enriching of our understanding of voters. Commercial matching technology has become quite sophisticated and solves many of the problems identified with earlier methods of political survey validation. Through validation, we can learn more about the nature of misreporting and, more importantly, about the nature of political participation.

Funding

Funds for this study came from the Dean of the Faculty of Arts and Sciences, Harvard University.

References

- Ansolabehere, Stephen, and Eitan Hersh. 2010. *The quality of voter registration records: A state-by-state analysis*. Caltech/MIT Voting Technology Project Report.
- . 2011a. Who Really Votes? In *Facing the challenge of democracy? Explorations in the analysis of public opinion and political participation*, eds. Benjamin Highton and Paul Sniderman. Princeton, NJ: Princeton University Press.
- . 2011b. *Gender by race: The massive participation gender gap in American politics*. Paper presented at the annual meeting of the American Political Science Association, Seattle, WA.
- . 2012. *Replication data for validation: What big data reveal about survey misreporting and the real electorate [dataset]*. IQSS Dataverse Network [Distributor] V1 [Version]. <http://hdl.handle.net/1902.1/18300> (accessed July 30, 2012).
- Ansolabehere, Stephen, Eitan Hersh, Alan Gerber, and David Doherty. 2010. *Voter registration list quality pilot studies: Report on detailed results and report on methodology*. Washington, DC: Pew Center on the States.
- 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.
- 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.
- Belli, Robert F., Sean E. Moore, and John VanHoewyk. 2006. An experimental comparison of question forms used to reduce voter overreporting. *Electoral Studies* 25(4):751–9.

- Berent, Matthew K., Jon A. Krosnick, and Arthur Lupia. 2011. *The quality of government records and "overestimation" of registration and turnout in surveys: Lessons from the 2008 ANES Panel Study's Registration and Turnout Validation Exercises*. American National Election Studies Working Paper No. nes012554.
- Bernstein, Robert, Anita Chadha, and Robert Montjoy. 2001. Overreporting voting: Why it happens and why it matters. *Public Opinion Quarterly* 65(1):22–44.
- Burden, Barry C. 2000. Voter turnout and the National Election Studies. *Political Analysis* 8(4):389–98.
- . 2003. Internal and external effects on the accuracy of NES turnout: Reply. *Political Analysis* 11(2):193–5.
- Burns, Nancy, Kay Lehman Schlozman, and Sidney Verba. 2001. *The private roots of public action*. Cambridge, MA: Harvard University Press.
- Campbell, James E. 2010. Explaining politics, not polls: Reexamining macropartisanship with recalibrated NES Data. *Public Opinion Quarterly* 74(4):616–42.
- Cassel, Carol A. 2003. Overreporting and electoral participation research. *American Politics Research* 31(1):81–92.
- . 2004. Voting records and validated voting studies. *Public Opinion Quarterly* 68(1):102–8.
- Citrin, Jack, Eric Schickler, and John Sides. 2003. What if everyone voted? Simulating the impact of increased turnout in senate elections. *American Journal of Political Science* 47(1):75–90.
- Denscombe, Martyn. 2006. Web-based questionnaires and the mode effect: An evaluation based on completion rates and data contents of near-identical questionnaires delivered in different modes. *Social Science Computer Review* 24(2):246–54.
- Deufel, Benjamin J., and Orit Kedar. 2010. Race and turnout in U.S. elections: Exposing hidden effects. *Public Opinion Quarterly* 74(2):286–318.
- Duff, Brian, Michael J. Hanmer, Won-Ho Park, and Ismail K. White. 2007. Good excuses: Understanding who votes with an improved turnout question. *Public Opinion Quarterly* 71(1):67–90.
- Fullerton, Andrew S., Jeffrey C. Dixon, and Casey Borch. 2007. Bringing registration into models of vote overreporting. *Public Opinion Quarterly* 71(4):649–60.
- Harbaugh, William T. 1996. If people vote because they like to, then why do so many of them lie? *Public Choice* 89(1):63–76.
- Highton, Benjamin, and Raymond E. Wolfinger. 2001. The political implications of higher turnout. *British Journal of Political Science* 31(01):179–223.
- Karp, Jeffrey A., and David Brockington. 2005. Social desirability and response validity: A comparative analysis of overreporting voter turnout in five countries. *Journal of Politics* 67(3):825–40.
- Katosh, John P., and Michael W. Traugott. 1981. The consequences of validated and self-reported voting measures. *Public Opinion Quarterly* 45(4):519–35.
- Katz, Jonathan N., and Gabriel Katz. 2010. Correcting for survey misreports using auxiliary information with an application to estimating turnout. *American Journal of Political Science* 54(3):815–35.
- Malhotra, Neil, and Jon A. Krosnick. 2007. The effect of survey mode and sampling on inferences about political attitudes and behavior: Comparing the 2000 and 2004 ANES to Internet surveys with nonprobability samples. *Political Analysis* 15(3):286–323.
- Martinez, Michael D. 2003. Comment on "voter turnout and the National Election Studies." *Political Analysis* 11(2):187–92.
- McDonald, Michael P. 2003. On the overreport bias of the National Election Study turnout rate. *Political Analysis* 11(2):180–6.
- . 2011. "Turnout, 1980–2010." United States Elections Project. http://elections.gmu.edu/voter_turnout.htm (accessed May 16, 2011).
- Presser, Stanley, Michael W. Traugott, and Santa Traugott. 1990. *Vote "over" reporting in surveys: The records or the respondents*. ANES Technical Report Series: Doc. nes010157. Ann Arbor, MI: American National Election Studies.
- Rosenstone, Steven J., and David C. Leege. 1994. An update on the National Election Studies. *PS: Political Science and Politics* 27(4):693–8.
- Rosenstone, Steven J., and John Mark 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(1):47–56.
- Silver, Brian D., Barbara A. Anderson, and Paul R. Abramson. 1986. Who overreports voting? *American Political Science Review* 80(2):613–24.
- Stanford University and the University of Michigan [producers and distributors]. 2010. American National Election Studies Time Series Cumulative Data File [dataset] version. <http://electionstudies.org/studypages/cdf/cdf.htm> (accessed June 24, 2010).
- Traugott, Michael W., Santa M. Traugott, and Stanley Presser. 1992. *Revalidation of self-reported vote*. ANES Technical Report Series: Doc. nes010160. Ann Arbor, MI: American National Election Studies.
- Verba, Sidney, Kay L. Schlozman, and Henry E. Brady. 1995. *Voice and equality: Civic voluntarism in American Politics*. Cambridge, MA: Harvard University Press.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who votes?* New Haven, CT: Yale University Press.