

American Politics Research

<http://apr.sagepub.com/>

Vote Self-Prediction Hardly Predicts Who Will Vote, And Is (Misleadingly) Unbiased

Todd Rogers and Masa Aida

American Politics Research published online 5 September 2013

DOI: 10.1177/1532673X13496453

The online version of this article can be found at:

<http://apr.sagepub.com/content/early/2013/09/05/1532673X13496453>

Published by:



<http://www.sagepublications.com>

Additional services and information for *American Politics Research* can be found at:

Email Alerts: <http://apr.sagepub.com/cgi/alerts>

Subscriptions: <http://apr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

>> [OnlineFirst Version of Record](#) - Sep 5, 2013

[What is This?](#)

Vote Self-Prediction Hardly Predicts Who Will Vote, and Is (Misleadingly) Unbiased

American Politics Research

XX(X) 1–26

© The Author(s) 2013

Reprints and permissions:

sagepub.com/journalsPermissions.nav

DOI: 10.1177/1532673X13496453

apr.sagepub.com



Todd Rogers^{1,2} and Masahiko Aida³

Abstract

Public opinion researchers, campaigns, and political scientists often rely on self-predicted vote to forecast turnout, allocate resources, and measure political engagement. Despite its importance, little research has examined the accuracy of self-predicted vote responses. Seven pre-election surveys with postelection vote validation from three elections ($N = 29,403$) reveal several patterns. First, many self-predicted voters do not actually vote (flake-out). Second, many self-predicted nonvoters do actually vote (flake-in). This is the first robust observation of flake-in. Third, actual voting is more accurately predicted by past voting (from voter file or recalled) than by self-predicted voting. Finally, self-predicted voters differ from actual voters demographically. Actual voters are more likely to be White (and not Black), older, and partisan than actual nonvoters (i.e., there is participatory bias), but self-predicted voters and self-predicted nonvoters do not differ much. Vote self-prediction is “biased” in that it misleadingly suggests that there is no participatory bias.

Keywords

political behavior, polling, public opinion, turnout, participatory bias

¹Harvard Kennedy School, Cambridge, MA, USA

²Analyst Institute, Washington, DC, USA

³Greenberg Quinlan Rosner Research, Washington, DC, USA

Corresponding Author:

Todd Rogers, Harvard Kennedy School, 79 John F. Kennedy Street, Cambridge, MA 02138, USA.

Email: Todd_Rogers@hks.harvard.edu

Of the 31 national public opinion polls released in the week before the 2012 U.S. Presidential election, only two accurately predicted the election outcome within two percentage points (Politico.com)—94% were off by more than two percentage points. Moreover, internal polling by Governor Romney's campaign that was disclosed after the election reflected substantial errors in forecasting which individuals would vote, and what the aggregate demographics of the voting electorate would look like. The question of who will vote is important to public opinion researchers, campaigns, and political scientists. Despite its importance, the validity and biases of self-predicted vote polling questions have been understudied. In this manuscript we examine the accuracy of self-predicted vote questions. We analyze a unique collection of seven pre-election surveys with postelection vote validation, involving a total of 29,403 interviews appended to validated postelection public records of turnout.¹ Five of these surveys were conducted during the 2008 U.S. Presidential Election, one survey was conducted during the 2009 New Jersey General Election, and one survey experiment was conducted during the Wisconsin Recall Elections in 2011.

Our findings speak to two lines of research. First, because we find that self-predicted vote is prone to several significant and consistent biases, our findings have implications for research that depends on accurate predictions of who will vote and who will not. Self-predicted vote is a commonly used dependent variable in political science as a proxy for political engagement.² For example, the original research on the effects of negative advertising on political engagement uses this question as its principal dependent variable to support the argument that “going negative” demotivates the electorate (Ansolabehere & Iyengar, 1995; Ansolabehere, Iyengar, & Simon, 1999; for a recent review, see Brooks, 2006). Similarly, research on televised incivility uses self-predicted vote to measure the consequences of combative television culture on political life (Mutz & Reeves, 2005). Work on how campaign activities influence political engagement uses self-predicted vote as a primary dependent variable (Hillygus, 2005). Self-predicted vote has also been used to show that citizens are more motivated to vote when they believe turnout will be relatively high as opposed to relatively low (Gerber & Rogers, 2009). And recent work on the impact of party affiliation on political behaviors found that inducing people to identify with a party made them more likely to self-predict that they would vote in from an upcoming election (Gerber, Huber, & Washington, 2010).

Making inferences about political engagement self-predicted vote is much more complicated and questionable in light of our results. To foreshadow our results we find both that a meaningful fraction of those who say they *will* vote do not (i.e., they “flake out”), and that that a large fraction of those who say

they will *not* vote, in fact, do (i.e., they “flake in”). We also find that *frequent* voters are more accurate predicting that they *will* vote, whereas *infrequent* voters are more accurate predicting that they will *not* vote. That is, respondents are more accurate predicting that they will behave consistently with their past behavior than predicting that they will behave inconsistently with their past behavior. Finally, using a survey experiment we rule out two possible memory-based explanations for why self-predicted vote is so inaccurate.

Second, our findings relate to research on how voters differ from nonvoters, a phenomenon called participatory bias (Brady, Verba, & Schlozman, 1995; Rosenstone & Hansen, 1993; Wolfinger & Rosenstone, 1980). To date there have been two kinds of participatory bias reported. The first, what we call “retrospective participatory bias,” compares voters and nonvoters based on postelection self-reported vote. The second, what we call “actual participatory bias,” compares voters and nonvoters based on publicly available administrative records of who votes (Ansolabehere & Hersh, 2011). There are two primary features of participatory bias. First, voters are more likely to be White, older, and partisan (i.e., registered as Democrat or Republican as opposed to no party affiliation) and less likely to be Black. Second, despite these demographic differences, voters and nonvoters are surprisingly similar in terms of policy attitudes and preferences. The retrospective participatory bias and the actual participatory bias literatures have found similar patterns for both findings, though the retrospective participatory bias measure exaggerates the demographic differences between voters and nonvoters relative to actual participatory bias measure. This difference is explained by parallel research on how accurate respondents are when they self-report whether they voted in postelection surveys. That work finds that respondents who are White, older, and partisan are most likely to overreport that they voted when they actually did not (for review, see Silver, Abramson, & Anderson, 1986b; Tourangeau, Rips, & Rasinski, 2000). Putting together research on overreporting and participatory bias, the attributes of respondents who are most likely to report that they voted when they did not are also the attributes that are most exaggerated when comparing retrospective participatory bias with actual participatory bias.

The current work examines participatory bias using a measure of participation that has not before been included in the scholarly discussion of participatory bias: pre-election surveys that solicit respondents’ self-predictions about whether they will vote. We call this “prospective participatory bias.” Of course, actual participatory bias offers the authoritative perspective on how representative voters are of nonvoters. That said, there are a range of important implications for the differences between those who self-predict that they will and will not vote, those who actually do and do not vote, and those who

report that they did and did not vote. Specifically, survey researchers, campaigns, and political scientists use pre-election self-predicted vote for a range of other uses—from developing campaign strategy for campaign targeting and resource allocation, to predicting who will win an election and by how much, to understanding the factors that influence political engagement. These and other uses of self-predicted vote could be improved by understanding the degree to which prospective participatory bias is related to bias in the electorate.

The participatory bias literature finds that self-reported and actual voters are more likely to be White, older, and partisan, and less likely to be Black than nonvoters. Our results differ substantially from that: We find that self-predicted voters are demographically similar to self-predicted nonvoters. Comparing self-predicted voters with actual voters, self-predicted voters are less likely to be White, older, and partisan, and more likely to be Black than actual voters. Relying on self-predicted vote to forecast who will vote leads to a peculiar “bias”: one would *mistakenly* forecast that voters will be demographically similar to nonvoters. The reason actual voters differ from self-predicted voters is that Black, younger, and nonpartisan respondents are more likely to flake out than White, older, and partisan respondents.

Contribution and Hypotheses

Several papers have addressed the accuracy of self-predicted vote, almost always focusing on identifying the optimal strategy for predicting turnout. A review of this research begins with Perry (1960, 1979) who wrestled with this problem extensively for The Gallup Organization. Perry developed several versions of a multiple-question battery primarily intended to identify who would actually vote versus those who were merely saying they would vote when, in fact, they would not. Several research teams have developed other multiple-question batteries using the most recent American National Election Studies (ANES) validated vote data from 1980, 1984, and 1988 (Freedman & Goldstein, 1996; Murray, Riley, & Scime, 2009). The most recent election for which we were able to find an analysis of the inaccuracies of self-predicted vote was the 1999 Philadelphia Municipal election. Dimock, Keeter, Schulman, and Miller (2001) aimed to improve on Perry’s multiple-question batteries by considering new questions and different model specifications. Those authors matched the interviews of 1,694 respondents to the municipal voter file and were thus able to determine if a given respondent voted or not. All of these datasets are small compared with the datasets used in this manuscript. In addition they are from elections held one to three decades ago, as opposed to the recent elections from which the current datasets come. Those

multiquestion batteries were developed primarily to reduce the bias toward offering responses that are perceived to be socially desirable—this is what we call “flake-out.” Such a bias is common in survey research, especially when there are no special strategies used to elicit accurate responses like incentives or list experiments (see Tourangeau et al., 2000). This leads to our first hypothesis:

Hypothesis 1: Some respondents will self-predict that they *will* vote, but they actually *will not* vote (i.e., they will flake out).

Unfortunately, because the vote validated ANES datasets are small they include very few respondents who self-predict not voting. However, the Dimock et al. (2001) study does include a number of self-predicted nonvoters. That survey included a self-predicted vote question using a 10-point scale, with 10 being the most likely to vote and 1 being the least likely to vote. Of the 1,694 respondents with validated vote, 7% offered an answer between 1 and 6, and 39% of them actually voted.³ That such a substantial fraction of self-predicted nonvoters, in fact, did vote is striking. To date, no datasets with validated vote have included many self-predicted nonvoters. This likely explains the lack of previous research on how common such an error is and why it might arise.

There are a range of reasons why self-predicted nonvoters might actually vote. First, respondents may strategically offer a self-prediction that they will not vote because they believe such a response will result in termination of the interview. Second, respondents may fail to anticipate the campaign and social forces leading up to Election Day. These unanticipated mobilizing forces may induce respondents to vote. For example, campaign get-out-the-vote activity may induce self-predicted nonvoters to actually vote (Green & Gerber, 2008). Or, self-predicted nonvoters may not anticipate that many of their friends and family will vote and therefore they could fail to anticipate the resulting social influence (Gerber & Rogers, 2009). Third, respondents may construe a question about whether they will vote as an opportunity to express disapproval for the political process, rather than as a genuine question about future behavior. This would be consistent with research suggesting that many voters are disaffected with politics, and that this feeling is unrelated to political participation tendencies (Pinkleton, Austin, & Fortman, 1998). Finally, self-predictions that respondents will not vote may be genuine errors in behavioral prediction. Research in cognitive and social psychology has shown that respondents’ sincere self-predictions of their future behavior are often only weakly related to their actual future behavior (see Theory of Planned Behavior: Ajzen, 1991). It is

conceivable that respondents simply lack the ability to accurately forecast their future behavior in the domain of voting. All of these rationales combine to lead to our second hypothesis:

Hypothesis 2: Some respondents will self-predict that they will not vote, but they actually will vote (i.e., they will flake in).

While we hypothesize that there will be substantial self-prediction errors, we also hypothesize that respondents' past voting behavior will strongly predict whether respondents vote, independent of their self-prediction. This is consistent with psychological research that has consistently shown that one of the best predictors of future behavior is past behavior (Ouellette & Wood, 1998; Triandis, 1977). If our results are consistent with this hypothesis, they would suggest a strategy for more accurately identifying which respondents will vote and which will not. Our third hypothesis is:

Hypothesis 3a: Respondent voting behavior will resemble their past vote histories.

Hypothesis 3b: Respondent past vote histories will be better predictors of whether they will vote than their self-predictions.

So far we have hypothesized that respondents' self-predictions about whether they will vote will be inaccurate (H1 and H2) and that respondents' vote histories will be more accurate predictors of whether they will vote than respondents' self-predictions (H3b). This suggests that if respondents can accurately recall their past vote history they can improve the accuracy of their self-predictions. We hypothesize that respondents will have introspective access to their past vote histories—though when asked to self-predict whether they will vote, respondents will tend to not automatically access it. However, we hypothesize that making respondents' vote histories cognitively accessible at the time of self-prediction will increase the degree to which that information will be incorporated into respondents' self-predictions, making self-predictions more accurate. This hypothesis is consistent with priming literature that suggests that attitudes are constructed based on available mental contents, and therefore treatments that affect respondents' mental contents can change their attitudes (see Payne, Bettman, Schkade, Schwarz, & Gregory, 2000). Put succinctly, we hypothesize:

Hypothesis 4a: Respondents can accurately recall their vote histories.

Hypothesis 4b: Respondents self-predictions about whether they will vote will be more accurate when respondents' vote histories are cognitively accessible.

Our next hypothesis predicts simply that prospective participatory bias will differ from actual participatory bias. Prospective participatory bias may emerge because some respondents are prone to saying that they will vote, but actually do not vote (“flake-out”), or say they will not vote, but actually do vote (“flake-in”). Prospective participatory bias could resemble actual participatory bias if the attributes of being White, older, and partisan are not associated with being likely to flake out or being likely to flake in; that is, if the accuracy of self-prediction is unrelated to those attributes.

On the other hand, prospective participatory bias may show the same exaggerated patterns of bias as retrospective participatory bias. This may occur if the attributes of being White, older, and partisan are positively associated with being likely to flake out (self-predicting that one will vote, but not actually voting), or negatively associated with being likely to flake in (self-predicting that one will not vote, but actually voting). The logic of overreporting research supports the “flake-out” mechanism because these attributes are positively associated with offering the socially desirable self-prediction (“yes, I will vote”) even when they will not actually vote. At the same time, the “flake-in” mechanism may be supported by the possibility that these attributes are positively associated with political sophistication, making these respondents more likely to use the self-prediction question as an opportunity to express dissatisfaction with politics, or more likely to answer the question strategically in an effort to terminate the survey.

Finally, prospective participatory bias could show even less bias than retrospective participatory bias and actual participatory bias. This could occur if the attributes of being White, older, and partisan are negatively associated with being likely to flake out, or positively associated with being likely to flake in. As these attributes are associated with political engagement (Brady et al., 1995), the Theory of Planned Behavior could support the flake-out mechanism (Ajzen, 1991). It argues that the more people care about something the less likely they are to not follow through (i.e., flake-out). At the same time, the flake-in mechanism could arise if campaigns target with get-out-the-vote activity people who are not highly likely to vote.

Because of these competing rationales, we offer our fifth hypothesis as nondirectional:

Hypothesis 5: Prospective participatory bias will differ from actual participatory bias and/or retrospective participatory bias.

Surveys

In this manuscript we examine seven pre-election surveys collected during three elections, each of which includes postelection vote validation. These

surveys differ along a range of dimensions. We will briefly describe the five surveys collected during the 2008 General Election, then the survey collected during the 2009 New Jersey General Election, and finally the survey collected during the 2011 Wisconsin Recall Election.

2008 General Election Surveys

We include five surveys from this election that use two types of sampling frames. All surveys were conducted to produce estimates of beliefs and behaviors of a representative sample of American likely voters. The first sampling frame involves four surveys conducted over 3 months using registration-based sampling based on voter files provided by Catalist, LLC (Green & Gerber, 2006). The sample frame was restricted to citizens with valid landline telephone numbers. The second sampling frame involves one survey that used random digit dial (RDD) based on valid landline phone numbers provided by Survey Sampling International. These surveys were conducted by Greenberg Quinlan Rosner Research, using the same call center, with the same supervising staff and training procedure. The samples include respondents from all 50 states and the District of Columbia, with data collection stratified by region. Table S1 provides details about each of the five surveys in this dataset.

All interviews included the following self-predicted vote question, "What are the chances of you voting in the election for president in November? Are you almost certain to vote, will you probably vote, are the chances 50-50, or don't you think you will vote?" If respondents answered "50-50," "will not vote" or volunteered that they "do not know," the interview was terminated. If respondents answered "almost certain to vote" or "will probably vote" the interviews continued, and generally lasted around 20 min. For calls to the RDD sample, respondents were asked their age, gender, and first name. For calls to the voter file sample, interviewers confirmed the name of the respondent. The proportion of respondents for whom we have data who refused to answer the vote self-prediction question is quite small (< 1%). This group is excluded from analyses.

After the election, surveys were merged with the national voter file provided by Catalist, allowing us to study 2008 turnout and past vote history.⁴ Some states do not disclose individual vote histories, and some states did not release vote histories in time to be included in the final dataset transferred to us for these analyses; therefore, interviews from 39 states were included in the present analyses.⁵ For interviews in the single survey that used a random-digit dial sampling frame, respondents were matched to the voter file within 2 months of completing the interview using the phone number dialed, the

respondent's gender as reported by the interviewer, and the respondent's self-reported age. Of these interviews, 54% matched the voter file using these criteria. For interviews completed calling from the voter file, respondents were matched after the election to the voter file using the name on the file from which the call was originally made, the phone number dialed, the respondent's gender as reported by the interviewer, and the respondent's self-reported age. Of these interviews, 93% matched the voter file using these criteria. Of the 12,979 interviews that included the self-predicted vote question, we matched 11,025 interviews to the 2008 General Election vote history. All analyses reflect only the data associated with the interviews that matched the voter file. Of these respondents, 53.5% are female and the average age is 57.

2009 NJ General Election Survey

This survey was not conducted with the current research in mind and was instead conducted on behalf of a 501 c(4) organization that supports non-partisan voter registration and get-out-the-vote efforts. Interviews were conducted during the 5 days before the election. Participants were randomly selected from a list of registered voters provided by Catalist, LLC, with unique, valid phone numbers who met three additional criteria: (a) were African American or Hispanic, (b) voted in the 2008 General Election, (c) did not vote in the 2006 General Election. Those who had already voted in the target election by the time of the interviews were also excluded.

Callers from a paid phone center introduced themselves and stated that the purpose of the call was to encourage the respondents to vote. They then asked the respondent "Do you plan to [vote/be a voter] on Tuesday?" There were two variants of the vote self-prediction question as part of an unrelated field experiment. The wording of this question did not affect the results reported in the present analysis, and so all analyses report the combined data. The response options were simply "Yes" or "No." A total of 13,278 responded to the vote prediction question. The calls continued as part of the unrelated, unpublished experiment. It found that emphasizing the voter identity as opposed to the fact that one can vote increased turnout by 1.2 percentage points (see related Bryan, Walton, Rogers, & Dweck, 2011). After the election, the surveys were merged with the voter file provided by Catalist, LLC, allowing us to study 2009 turnout and past vote history. As the initial participant list was drawn from the Catalist, LLC, voter file and callers were instructed to talk only to the targeted person, and all respondents were matched. 51.5% of respondents are female and the average age is 42.7.

2011 Wisconsin State Senate Recall Elections Survey

This survey was not conducted with the current research in mind and was instead conducted on behalf of the AFL-CIO, a leading labor union. The objective of the interviews was to help the AFL-CIO build microtargeting models to identify which citizens would support which candidates in the elections to be held in July and August. Interviews were conducted during July and August in 2011. Registered voters with unique and valid phone numbers were randomly selected from a voter file provided by Catalist, LLC.

This survey included questions that allow us to examine H4a and H4b. Callers from a paid phone center introduced themselves and stated that the purpose of the call was to ask questions about the upcoming election. All respondents were asked two key questions and were randomly assigned to which one they answered first. Those assigned to *Recall-First* were asked to recall whether they had voted in a Supreme Court election held the previous April. This was the most recent previous election that was expected to have a turnout level similar to that of the forthcoming recall election. They were asked “*There was an election for Supreme Court held on April 5 of this year. Were you able to make it to the polls and vote in this election?*” The response options were “voted,” “did not vote,” or “don’t know/refused.” Those assigned to *Self-Prediction-First* were asked to self-predict if they would vote in the upcoming election. They were asked “*As you may have heard, there will be an election to recall [INCUMBENT PARTY] state senator [INCUMBENT NAME]. Do you intend to vote in this upcoming election?*” The coded response options were “definitely,” “probably,” “50/50,” “will not vote,” and “don’t know/refused.”

A total of 5,100 respondents answered both questions. The interview completion rate was similar across conditions, and respondents across conditions did not significantly differ on demographic variables available on the file such as gender, age, percent Democrat, percent Republican, or turnout in April Supreme Court vote history (all $ps > .10$). The current analyses include only respondents who answered both questions. After the election, surveys were merged with the voter file provided by Catalist, LLC, allowing us to study 2011 turnout and past vote history. As the initial participant list was drawn from the Catalist, LLC, voter file, and callers were instructed to talk only to the targeted person, and all respondents were matched. A total of 53.8% of respondents are female and the average age is 60.4.

Results

Combined, H1 and H2 predict two forms of self-prediction inaccuracy: some respondents would self-predict that they *would* vote, but they actually *would*

Table 1. Self-Prediction Accuracy Across Elections.

		2008 general election			2009 New Jersey general election			2011 Wisconsin recall election		
		Actually vote?			Actually vote?			Actually vote?		
		Yes (%)	No (%)	Total N	Yes (%)	No (%)	Total N	Yes (%)	No (%)	Total N
Self-predicted likelihood to vote	[Yes]	86.7	13.3	9,272	45.8	54.2	8,798	82.7	17.3	4,021
	Probably	74.4	25.6	1,004				59.6	40.4	282
	50–50	66.7	33.3	417				45.3	54.7	232
	[No]	54.7	45.3	201	29.3	70.7	3,149	39.5	60.5	473
	[Don't know]	76.3	23.7	131	42.4	57.6	1,318	51.1	48.9	92
Total		84.1	15.9	11,025	41.5	58.5	13,265	24.9	75.1	5,100

Note. We cannot prove that a respondent did not vote. We report here that the Secretary of State of the state in which a given respondent lives does not indicate the respondent had cast a vote. In the 2008 General Election, the question asked was: What are the chances of you voting in the election for president in November: are you almost certain to vote, will you probably vote, are the chances 50-50, or don't you think you will vote? (We were not able to collect information that allows us to assess the American Association of Public Opinion Researchers (AAPOR) response rate 2). In the 2009 New Jersey General Election (AAPOR response rate 2: 30.2%), the question asked was: Do you plan to [be a voter/vote] on Tuesday? In the 2011 Wisconsin Recall Election (AAPOR response rate 2: 9.0%), the question asked was: As you may have heard, there will be an election to recall [INC_PART] state senator [INC_NAME]. Do you intend to vote in this upcoming election? INC = incumbent.

not vote (flake-out); while some respondents would self-predict that they would not vote, but they actually would vote (flake-in). Table 1 shows that results from all three elections are consistent with H1 and H2. Consistent with H1, a sizable fraction of those who self-predicted that they would vote mispredicted and did not actually vote (flake-out). In the 2008 General Election surveys this was the case for 13.3% of self-predicted voters; in the 2009 New Jersey General Election survey this was the case for 54.2% of self-predicted voters; and, in the 2011 Wisconsin Recall Election survey this was the case for 17.3% of self-predicted voters. This is consistent with social desirability bias research, according to which people may have knowingly self-predicted that they would vote when they knew they would not, but this could also result from unanticipated obstacles preventing respondents from voting despite their otherwise genuine self-prediction.

Consistent with H2, a sizable fraction of those who self-predicted that they would not vote mispredicted and actually did vote (flake-in). In the 2008 General Election surveys this was the case for 54.2% of self-predicted nonvoters; in the 2009 New Jersey General Election survey this was the case for 29.3% of self-predicted nonvoters; and, in the 2011 Wisconsin Recall Election survey this was the case for 39.5% of self-predicted nonvoters. The rate of flake-in is surprisingly large. In fact, in the case of the 2008 General Election

surveys and the 2011 Wisconsin Recall Election survey this fraction was larger than the fraction who made the more widely known misprediction: self-predicting that they would vote, but actually not voting, $\chi^2(1, N = 9,473) = 166.3, p < .001$, and $\chi^2(1, N = 4,494) = 461.6, p < .001$.

Hypothesis 3a predicted that respondents' voting behavior would be consistent with their past vote histories. Table 2 shows that results from all three elections are consistent with this hypothesis. These tables show that respondents are more accurate when their self-predictions are consistent with their past voting behavior. Respondents in the 2008 General Election surveys who had voted in the past two General Elections were more accurate when predicting that they "were almost certain to vote" or would vote (93% of those predicting that they "were almost certain to vote" actually did vote; 93% accuracy because accurate prediction for this response option entailed actually casting a vote) than that they "will not vote" (76% of those predicting that they "will not vote" actually did vote; 24% accuracy because accurate prediction for this response option entailed not casting a vote). Similarly, respondents who had voted in neither of the prior two General Elections were more accurate at predicting that they "will not vote" (30% of those predicting that they "will not vote" actually did vote; 70% accuracy because accurate prediction for this response option entailed not casting a vote) than that they were "almost certain to vote" (62% of those predicting that they were "almost certain to vote" actually did vote; 62% accuracy because accurate prediction for this response option entailed actually casting a vote). A logistic regression shows that the interaction between vote history and predictions on accuracy is statistically significant (odds ratio = .12, $p < .001$).

Respondents in the 2009 New Jersey General Election survey who had voted in both of the past two New Jersey General Elections were more accurate when predicting that they would vote (80% of those predicting that they would vote actually did vote; 80% accuracy) than that they would not vote (66% of those predicting that they would not vote actually did vote; 34% accuracy). Similarly, respondents who had voted in none of the past two New Jersey General Elections were more accurate at predicting that they would not vote (24% of those predicting that they would not vote actually did vote; 76% accuracy) than that they would vote (39% of those predicting that they would vote actually did vote; 39% accuracy). A logistic regression shows that the interaction between vote history and predictions on accuracy is statistically significant (odds ratio = 6.71, $p < .001$).

The 2011 Wisconsin Recall Election survey occurred in a highly unusual recall election context. To identify the two most comparable elections we selected the previous two elections that had similar overall turnout. These

Table 2. Actual Turnout by Self-Predicted Turnout by Vote History.

Self-predicted likelihood to vote	Votes cast in previous two general elections											
	2008 general election			2009 New Jersey general election			2011 Wisconsin recall election			2011 Wisconsin recall election		
	None	One	Two	None	One	Two	None	One	Two	None	One	Two
[Yes]	64.8% (n = 863)	80.1% (n = 1,890)	93.5% (n = 6,017)	39.8% (n = 5,874)	60.3% (n = 1,747)	79.7% (n = 518)	56.3% (n = 435)	79.7% (n = 1,121)	91.6% (n = 2,257)			
Probably	43.8% (n = 160)	72.0% (n = 246)	88.4% (n = 510)				28.1% (n = 64)	61.7% (n = 107)	85.5% (n = 83)			
50-50	43.5% (n = 92)	71.5% (n = 123)	81.3% (n = 166)				28.0% (n = 75)	45.6% (n = 79)	73.3% (n = 60)			
[No]	31.5% (n = 54)	54.2% (n = 59)	75.7% (n = 70)	24.2% (n = 2,284)	48.9% (n = 509)	66.4% (n = 122)	16.0% (n = 163)	45.5% (n = 134)	81.8% (n = 110)			
[Don't know]	48.1% (n = 27)	71.0% (n = 31)	91.2% (n = 68)	36.1% (n = 894)	62.5% (n = 253)	76.0% (n = 75)	28.6% (n = 21)	65.6% (n = 32)	75.0% (n = 24)			

Note. We cannot prove that a respondent did not vote. We report here that the Secretary of State of the state in which a given respondent lives does not indicate the respondent had cast a vote. In the 2008 General Election, the question asked was: 'What are the chances of you voting in the election for president in November: are you almost certain to vote, will you probably vote, are the chances 50-50, or don't you think you will vote?' (We were not able to collect information that allows us to assess the American Association of Public Opinion Researchers [AAPOR] response rate 2.) In the 2009 New Jersey General Election (AAPOR response rate 2: 30.2%), the question asked was: 'Do you plan to [be a voter/vote] on Tuesday?' In the 2011 Wisconsin Recall Election (AAPOR response rate 2: 9.0%), the question asked was: 'As you may have heard, there will be an election to recall [INC_PART] state senator, [INC_NAME]. Do you intend to vote in this upcoming election?' INC = incumbent.

were the April 2011 Supreme Court election and the February 2008 Presidential Primary. Table 2 reports the turnout rate among respondents in the present study who had voted in zero, one, or two of those elections.⁶ Respondents in this election who had voted in both of the most comparable past elections in Wisconsin were more accurate when predicting that they would vote (92% of those predicting that they would vote actually did vote; 92% accuracy) than that they would not vote (81% of those predicting that they would not vote actually did vote; 19% accuracy). Similarly, respondents who had voted in zero of the most comparable past elections in Wisconsin were more accurate at predicting that they would not vote (15% of those predicting that they would not vote actually did vote; 85% accuracy) than that they would vote (52% of those predicting that they would vote actually did vote; 52% accuracy). A logistic regression shows that the interaction between vote history and predictions on accuracy is statistically significant (odds ratio = .063, $p < .001$).

Figure 1 shows graphically the accuracy depicted in Table 2 by plotting the accuracy of respondents' turnout self-predictions crossed with their recent vote history for each election.

H3b predicted that actual voting would be more accurately predicted by past voting than by self-predicted voting. This was true across all three elections. To compare the explanatory power of the turnout self-prediction question to that of past vote history, we used a series of logistic regression models all of which controlled for age and gender. For the 2008 General Election surveys, including only self-predicted likelihood to vote, explained 4.0% of variance (pseudo R^2) whereas including past vote history alone explained 10.8% of variance. The model with self-prediction and past vote history explained 12.3% of variance. For the 2009 New Jersey General Election survey, including only self-predicted likelihood to vote explained 3.5% of variance (pseudo R^2) whereas including past vote history alone explained 5.5% of variance. The model with self-prediction and past vote history explained 6.9% of variance. For the 2011 Wisconsin Recall Election survey, including only self-predicted likelihood explained 13.6% of variance (pseudo R^2) whereas including only past vote history explained 18.1% of variance. The model with self-prediction and past vote history explained 22.1% of the variance. As a point of comparison, an ordinary least squares model with age and gender as covariates that only included respondents' self-reported recollection of whether they voted in the April election explained 13.8% of variance—It was slightly better at predicting future turnout than self-prediction of future turnout.

Across elections, self-prediction explained less variance than past vote history. Moreover, the models that include self-prediction and past vote




Figure 1. Respondents are more accurate when predicting past-consistent behavior, across studies.

Note. Self-prediction accuracy refers to the percentage of respondents who actually voted after self-predicting they would vote (“self-predicted voters” is solid, blue line) and who did not vote after self-predicting they would not vote (“self-predicted nonvoters” is dashed, red line). 2008 General election self-predicted voters offered response option “Almost certain to vote” while self-predicted nonvoters offered response option “Will not vote.” 2009 New Jersey General Election self-predicted voters offered response option “Yes, I will vote” while self-predicted nonvoters offered response option “No, I will not vote.” 2011 Wisconsin Recall Election self-predicted voters offered response option “Definitely voting” while self-predicted nonvoters offered response option “Will not vote.”

history do not explain much more variance than models that just include past vote history. Furthermore, even self-reported recollection of past voter history was a better predictor of voting in a specific upcoming election than self-prediction of whether one would vote in that specific election.

H4a predicted that respondents would have introspective access to accurate recollections about whether they voted in a past election. The 2011 Wisconsin Recall Election survey provided an opportunity to test this hypothesis. The vast majority of respondents were correct in their recollections:

80.1% ($SE = 0.006$) of respondents correctly recalled whether they had voted in the April 2011 Supreme Court elections. Among those who said that they did vote in the April 2011 Supreme Court ($n = 3,743$), 79.3% ($SE = 0.007$) actually did vote (20.7% did not). Among those who said that they did not vote in the April 2011 Supreme Court ($n = 1,257$), 82.1% ($SE = 0.010$) actually did not vote (17.9% did vote). Among those who say they did not know or refused to answer the question ($n = 92$), 41.3% ($SE = 0.052$) actually did vote (58.7% did not). Accuracy was not affected by whether the respondent was in *Recall-First* or *Predict-First* ($\chi^2 = .01, p > .9$). This suggests that respondents have access to reasonably accurate information regarding their past vote history, despite the motivational biases people have to over-report their vote history (Silver, Anderson, & Abramson, 1986a; Tourangeau et al., 2000), as well as the cognitive challenges of accurately remembering past experiences (Schacter, 1999).

H4b predicts that one of the reasons respondents do not make accurate self-predictions is that relevant past behavior is not cognitively available when they are formulating their self-predictions. The Wisconsin 2011 Recall Election survey experiment helps to assess this by asking respondents to recall their past turnout before predicting their future turnout (as opposed to recalling it after they have already predicted their future turnout. H4b predicts that accuracy should be greater for respondents in *Recall-First* than for respondents in *Recall-Second*. As the turnout prediction question had five response options, we used the response option indicating the greatest degree of certainty that one would vote ("I will definitely vote") as a prediction that one would vote. We used the response option indicating the greatest degree of certainty that one would not vote ("I will not vote") as a prediction that one would not vote. The relationship between all other prediction options and actual turnout is not incorporated into our criterion for prediction accuracy. Among those in *Recall-First* condition, 80.4% correctly predicted whether they would vote. Among those in the *Prediction-First* condition, 80.3% correctly predicted whether they would vote.

As question order did not affect accuracy of past vote recall, we are able to compare the prediction accuracy of those for whom accurate past behavior from a recent similar election was made salient before they predicted their voting likelihood (those who accurately recalled their past voting behavior in the *Recall-First* condition) to those for whom that information was not made salient before predicting whether they will vote (those who accurately recalled their past voting behavior in the *Predict-First* condition). Among those who accurately recalled their past vote history, making past vote history salient before predicting future turnout did not improve respondents' prediction accuracy, $\chi^2(1) = 0.020, p = .888$. Among all respondents,

Table 3. Actual Participatory Bias vs. Prospective Participatory Bias.

		% White	% Black	Average age	% No Party Affiliation	% Partisan
2008 general election	Actual voters	87.3	6.4	57.6	41.7	58.3
		0.3	0.3	0.487	0.5	0.5
	Actual nonvoters	82.7	10.6	53.6	58.1	41.9
		0.9	0.7	0.18	1.2	1.2
	Self-predicted voters	86.8	7.0	56.75	43.6	56.4
		0.4	0.3	0.18	0.5	0.5
	Self-predicted nonvoters	86.9	7.3	56.34	44.5	55.5
		2.4	1.9	1.475	3.6	3.6
	Actual participatory bias	4.6	-4.2	4	-16.4	16.4
	Prospective participatory bias	-0.1	-0.3	0.41	-0.9	0.9
t-statistic (actual vs. prospective bias)	21.6**	18.0**	13.0**	37.9**	37.9**	

Note. Note that White and Black do not add to 100% because some respondents do not have a race listed. *Signifies $p < .05$. **signifies $p < .01$.

recalling their vote history before predicting their future turnout slightly (but inconsequentially) changed respondents' turnout predictions, $\chi^2(4, 5100) = 12.49, p = .014$. In the *Recall-First* condition, 1.0 percentage point more respondents predicted they would definitely vote (79.3%) compared with those in the *Predict-First* (78.3%). When dichotomizing the respondents into those who predicted they would definitely vote and all others, we find recalling vote history does not affect prediction accuracy, $\chi^2(1, 5100) = 0.01, p = .90$. These findings are inconsistent with H4b.

H5 made the nondirectional prediction that prospective participatory bias would differ from the patterns of actual participatory bias and/or retrospective participatory bias. To analyze this hypothesis we study only the five surveys from the 2008 General Election and not the 2009 New Jersey General Election survey or the Wisconsin 2011 Recall Election survey. We restrict to these surveys because they are representative of eligible voters and questions about participatory bias call for a diverse sampling of respondents. Conversely, the sampling strategies for the two surveys from the other two elections are specifically biased toward targeted demographics, making interpretation about prospective participatory bias unclear. Table 3 shows the crosstabs of the five surveys from the 2008 General Election.

We replicate the Ansolabehere and Hersh (2011) findings regarding actual participatory bias: those who actually vote are more likely to be White, less

likely to be Black, more likely to be older, and more likely to be partisan (i.e., registered Democrat or Republican) than those who actually do not vote.

Despite the robustness of actual participatory bias, Table 3 shows that there is little prospective participatory bias: self-predicted voters very closely resemble eligible voters. In general, self-predicted vote creates a misleadingly unbiased estimate of who will vote relative to what actual voter records suggest.

The data allow us to explore some mechanisms of why self-predicted voters differ significantly from actual voters in terms of race, age, and partisanship. The reader will recall that we discussed two potential mechanisms for this. The first mechanism involves flake-out: respondents self-predict that they will vote, but they actually do not vote. Flake-out could explain the patterns we find if Black, younger, and nonpartisan respondents are more likely to flake out than White, older, and partisan respondents. The second mechanism involves flake-in: respondents self-predict that they will not vote, but they actually do vote. The flake-in mechanism could explain the patterns we find if Black, younger, and nonpartisan respondents are less likely to flake in than White, older, and partisan respondents.

To analyze this question we create for each attribute a four cell matrix in which we cross the percentage of respondents from the 2008 General Election surveys with a given attribute who flake in and the percentage of respondents with that attribute who flake out, as a percentage of the total respondents in a specific category (e.g., partisan, nonpartisan, etc.). For example, the matrix for partisan and nonpartisan respondents looks like the following: percentage of partisan respondents who flake in (3.4%), percentage of nonpartisan respondents who flake in (3.7%), percentage of partisan respondents who flake out (10.0%), and percentage of nonpartisan respondents who flake out (17.9%), $\chi^2(1) = 114.8, p = < .001$. The matrix for White and Black respondents looks like the following: percentage of White respondents who flake in (3.4%), percentage of Black respondents who flake in (4.3%), percentage of White respondents who flake out (12.9%), and percentage of Black respondents who flake out (20.8%), $\chi^2(1) = 33.8, p = < .001$. To analyze age we perform a median split, which means we break it into two roughly equally sized categories: younger (18-57) and older (57-101). The matrix for younger and older respondents looks like the following: percentage of younger respondents who flake in (3.2%), percentage of older respondents who flake in (3.9%), percentage of younger respondents who flake out (14.4%) and percentage of older respondents who flake out (10.8%), $\chi^2(1) = 29.4, p = < .001$. These analyses support the interpretation self-predicted voters are less demographically biased than actual voters because nonpartisan, Black, and younger respondents self-predict that they will vote, but they fail to cast ballots.

General Discussion

Self-predictions about whether a person will vote are surprisingly inaccurate. Many of those who self-predict that they will vote, in fact, do not vote (flake-out), and many of those who self-predict that they will not vote, in fact, do vote (flake-in). Flake-in has not been robustly measured before. These errors in self-prediction are significantly correlated with whether respondents predict that they will behave consistently with their own past behavior. As a predictor of actual voting, self-prediction is far inferior to past vote history—and this is not because past vote history is not accessible or salient to respondents when they generate their self-predictions. Moreover, self-predicted turnout is no better at predicting actual turnout than self-reported recall of turnout in a single past election. Finally, self-predicted voters differ quite substantially from actual voters in terms of demographics: actual voters are more likely to be White, less likely to be Black, more likely to be older, and more likely to be partisan than self-predicted voters. While actual voters tend to be higher socioeconomic status (SES) than nonvoters (i.e., there is a “participatory bias”), self-predicted voters resemble nonvoters more so than actual voters in terms of these attributes. This difference emerges because Black, younger, and nonpartisan respondents tend to flake out, and not because White, older, and partisan respondents tend to flake in.

There are several implications of these findings. First, as reviewed above, these findings raise questions about the usefulness of self-predicted vote as a dependent variable in political science research. These findings do not altogether undermine the use of this dependent variable, but rather raise new questions. For example, given the inaccuracy of self-predicted vote highlighted in this manuscript, how should we interpret the meaning of a change in self-predicted vote, as is common in the literature described above (e.g., Ansolabehere & Iyengar, 1995; Hillygus, 2005; Mutz & Reeves, 2005; Gerber et al., 2010; Gerber & Rogers, 2009)? Does it signify an increased probability of voting, a fleeting change in motivation to vote, or possibly neither? If self-predicted vote questions are used in that research as a proxy for voter engagement/excitement/enthusiasm for an upcoming election, then perhaps questions specifically tailored to those constructs might be more useful. One implication of these findings for political science research is that analysis of data involving self-predicted vote should explicitly address past vote history—either by controlling for it in observational data or by stratifying by it in experimental data. This is because past vote history is a much more powerful predictor of future turnout than self-prediction is. Across all three studies, adding self-prediction to a predictive model that includes past vote history increases variance explained by less than a quarter, whereas

adding past vote history to a model that includes self-prediction more than doubles variance explained. Given that the three largest, on-going multiresearcher political survey research projects (Cooperative Congressional Election Study [CCES], General Social Survey [GSS], and ANES) include self-predicted vote questions, it is likely that future research will continue to use self-predicted vote as a dependent variable, which makes these considerations of continuing relevance.

Second, this research has implications for voter registration policy. Recall that a large proportion of eligible voters who predicted that they would not vote actually did vote. This suggests that scheduling the deadline for voter registration in advance of an election (as all but a small number of states currently do) may reduce overall turnout. Registration in advance of an election requires that citizens anticipate their interest for casting a vote in an election, and the results we report show that respondents' ability to do this is limited. This is consistent with recent research suggesting that Election Day registration increases overall turnout (Highton, 2004; Knack, 2001).

Third, this research has implications for political researchers and political campaigns. Political pollsters commonly screen who is eligible to participate in their surveys by including only those who report being "absolutely certain to vote" or "very likely going to vote." The present research suggests that this screen question is grossly inadequate for defining the "likely electorate."⁷ Moreover, it creates a systematically biased image of the likely electorate; it overestimates the degree to which Black, younger, and nonpartisan respondents will participate in an election. One approach we suggest is using a combination of self-prediction and voter file data. This is effectively using registration-based sampling and putting some weight on self-prediction (Green & Gerber, 2006). For example, the present data suggest that when calling from a list of registered voters a hybrid approach incorporating vote history and self-prediction could substantially increase accuracy when predicting who will vote and who will not vote. Accounting for other variables, such as demographic characteristics available on voter files, may make this strategy even more effective. In addition, our research has immediate implications for practitioners involved in campaign activities who use self-predicted vote as a screen for whether to target citizens with persuasive communications and/or get-out-the-vote contact.

Why do such a large proportion of respondents who self-predicted that they would not vote actually vote? The survey experiment conducted in the 2011 Wisconsin Recall Election tests and rules out two memory-based hypotheses. First, it rules out the hypothesis that respondents do not have access to accurate past vote history information because respondents were more than 80% accurate when recalling whether they voted in a specific past

election. And second, it rules out the hypothesis that past vote history is not salient when predictions are being formed by showing that making past vote history salient before self-predicting whether one would vote in an upcoming election did not improve accuracy. These findings either mean that the relevant past information is already salient when respondents are predicting their future voting behavior, which would be consistent with mounting evidence from psychology and neuroscience (see Schacter, Addis, & Buckner, 2007), or that, for this particular domain, past behavior is not incorporated into self-predictions of future behavior. With the current data we are unable to disentangle these two possible interpretations, but these results are consistent with behavioral decision research showing that people underuse relevant base-rate information when forecasting future behaviors (Kahneman & Lovallo, 1993). That work argues that people make forecasts from an “inside” perspective, believing that the current extenuating circumstances make information from an “outside” perspective (i.e., base-rate information) nondiagnostic.

A second possible explanation is that respondents are just making a forecasting error because they are so temporally distant from Election Day when they are making their self-predictions (e.g., perhaps because of a failure to anticipate the energy and excitement as Election Day neared). This would suggest that self-prediction should become more accurate as the prediction is made nearer to Election Day. The surveys from the 2008 General Election can help to answer this question. Because surveys were collected over the course of 5 months leading up to Election Day we can ask, “Were self-reported likelihoods of voting more accurate as the election approached?” One would suspect that as Election Day neared respondents might have had introspective access to information useful to predicting their future behavior. If respondents gained insight into their likely voting behavior as the election approached, we would hypothesize that errors in turnout self-predictions would decrease as Election Day approached. This is not confirmed, as there appears to be no trend in accuracy, $r(10,894) = .012$, $p = .21$. This suggests that the inaccuracy of self-reported vote is not caused by a failure to anticipate Election Day enthusiasm.

Third, it is conceivable that respondents expected that reporting that they would not vote would result in termination of the interview and so they strategically offered that response to politely get off the phone. If this were the case we might expect that in the 2008 General Election surveys, among the four response options available, there would be an especially high actual turnout rate among those who reported that they “will not vote” relative to those who reported that they had a “50-50 chance” of voting. This was not the case, and Table 2 shows that turnout rate increases virtually linearly with each level of response option.

Fourth, when respondents self-predicted that they would not vote it is possible that they failed to appreciate the social and informational experience that would arise on Election Day. By failing to account for the additional motivation aroused by the excitement of others, they may have underestimated the motivation they would feel when it came time to actually cast ballots. For example, when self-predicting whether they would vote days before Election Day, respondents may not have been aware that many of their friends and family were going to vote and therefore they may have failed to anticipate the resulting social influence (Gerber & Rogers, 2009). If this explanation is true, relatively high profile elections with high energy around Election Day and substantial spending at the end of the campaign should show higher rates of flake-in relative to lower energy and lower profile elections. Consistent with this hypothesis, we found that the flake-in rate among respondents in the 2008 General Election surveys (54.7%) was greater than the flake-in rate among respondents in the two lower salience elections: the 2009 New Jersey General Election (29.3%) and the 2011 Wisconsin Recall Election (39.5%), $t(1517) = 7.2, p < .001$ and $t(672) = 3.6, p < .001$, respectively. We do note, though, the limited validity of comparing just these three elections because they involved different voting likelihood questions, different sampling procedures, and were conducted at different times in their respective election cycles. We hope future research will test this hypothesis directly.

Fifth, given that the respondents who flaked-in were disproportionately regular voters, it is possible that these respondents offered answers meant to convey disaffection toward the political process rather than sincere belief that they would not vote. This would be consistent with research suggesting that many voters are disaffected with politics, and that this feeling is unrelated to political participation tendencies (Pinkleton et al., 1998). Future research could further examine this hypothesis in validated vote surveys by asking questions after the self-predicted vote question about respondents' feelings toward politics.

Sixth, it is possible that respondents genuinely offered their best predictions about whether they would vote and that they are simply poor predictors. The present data do not allow us to further assess this possibility, but it is conceivable and consistent with psychological research on the limits of behavioral prediction. Interestingly, the 2011 Wisconsin Recall Election survey shows that respondents have access to information that would be highly predictive of whether they would vote (i.e., their past behavior). For reasons we do not understand, respondents seem to not correctly access or weigh that information when predicting whether they will vote. We look forward to further research on this topic as well.

Acknowledgments

We would like to thank Eitan Hersh, David Nickerson, Don Green, Mike Bailey, Ryan Enos, and Dan Hopkins for their helpful feedback, and Debby Kermer, Julia Kamin, and Carly Robinson for help preparing this manuscript. We would like to thank Greenberg Quinlan Rosner Research, Democracy Corps, AFL-CIO, Catalist, LLC, and an anonymous nonprofit for sharing their data with us.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. Recent research has explored the nuances, complexities, and challenges of matching randomly sampled interviews to administrative records of election participation from Secretaries of States (see Berent, Krosnick, & Lupia, 2011 for a detailed discussion of some of the challenges of validated vote matching). This is a difficult task. The present studies are not vulnerable to most of the problems that work encounters, however. The sampling frames for six of the seven surveys used in this research derived exclusively from pre-existing voter files. This means the respondents were selected based on being prematched to previous administrative records. Although this may challenge the degree to which our sampling frames are representative of the country at large (which is not central to our research aims), it dramatically increases the quality of the postelection vote validation. The voter files we use, provided by Catalist LLC, were developed and were maintained using administrative records from Secretaries of State. The current manuscript reflects one of a handful of recent research projects utilizing this new resource for political research, though undoubtedly others are soon to follow. (For a discussion of how the Catalist, LLC, voter file is managed and validated, see Ansolabehere & Hersh, 2010).
2. Psychologists also use questions involving self-predicted vote. For example, Fast, Gruenfeld, Sivanathan, and Galinsky (2009) use self-predicted vote as a dependent variable to show the effects of increasing perceived power on general engagement.
3. Mann (2003) finds a similar result but with a very small sample size: 27 respondents offer “less than 50-50” or “Definitely will not vote” of which 7 actually did vote (26%).
4. Recent research has used and validated the data reported in the Catalist voter file (see Ansolabehere & Hersh, 2012; Ansolabehere, Hersh, & Shepsle, 2012)

5. The 39 states included in the analysis are AK, AL, AR, CO, DE, FL, GA, IA, ID, IL, IN, KS, KY, LA, MA, MD, ME, MI, MN, MO, MS, MT, NC, NH, NJ, NM, NV, NY, OH, OK, OR, PA, RI, SD, TN, UT, VT, WA, and WI.
6. For the 2011 Wisconsin Recall Election survey, we show in Table 3 the actual turnout by vote history in the past two comparable elections because this is the format of the data shown in the 2008 General Election surveys and the 2009 New Jersey General Election survey. In the 2011 Wisconsin Recall Election survey experiment we only asked respondents to recall whether they voted in the single most recent similar election, the April Supreme Court election. The same pattern holds when examining these data where vote history includes only none and one votes cast. For example, among those who did vote in the April election, 90% of those who self-reported that they “definitely will vote” actually voted, while 76% of those who reported that they “will not vote” actually voted. Conversely, among those who did not vote in the April election, 61% of those who self-reported that they “definitely will vote” actually voted, while 20% of those who reported that they “will not vote” actually voted.
7. As these surveys do not include interviews from those who flake in, we are not able to determine here how different the preferences of actual voters were from those who self-predicted that they would vote.

References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50, 179-211.
- Ansolabehere, S., & Hersh, E. (2010). *The quality of voter registration records: A state-by-state analysis*. Cambridge, MA: Department of Government, Harvard University.
- Ansolabehere, S., & Hersh, E. (2011). Who really votes? In P. M. Sniderman & B. Highton (Eds.), *Facing the challenge of democracy: Explorations in the analysis of public opinion and political participation* (pp. 267-291). Princeton, NJ: Princeton University Press.
- Ansolabehere, S., & Hersh, E. (2012). Validation: What big data reveal about survey misreporting and the real electorate. *Political Analysis*, 20, 437-459.
- Ansolabehere, S., Hersh, E., & Shepsle, K. (2012). Movers, stayers, and registration: Why age is correlated with registration in the U.S. *Quarterly Journal of Political Science*, 7, 1-31.
- Ansolabehere, S., & Iyengar, S. (1995). *Going negative: How political advertisements shrink and polarize the electorate*. New York, NY: Free Press.
- Ansolabehere, S., Iyengar, S., & Simon, A. (1999). Replicating experiments using aggregate and survey data: The case of negative advertising and turnout. *American Political Science Review*, 93, 901-909.
- Berent, M. K., Krosnick, J. A., & Lupia, A. (2011). *The quality of government records and “over-estimation” of registration and turnout in surveys: Lessons from the 2008 ANES Panel Study’s registration and turnout validation exercises* (Working

- Paper No. nes012554). Ann Arbor, MI: American National Election Studies. Retrieved from <http://www.electionstudies.org/resources/papers/nes012554.pdf>
- Brady, H. E., Verba, S., & Scholzman, K. L. (1995). Beyond SES: A resource model of political participation. *American Political Science Review*, *89*, 271-294.
- Brooks, D. J. (2006). The resilient voter: Moving toward closure in the debate over negative campaigning and turnout. *The Journal of Politics*, *68*, 684-696.
- Bryan, C. J., Walton, G. M., Rogers, T., & Dweck, C. S. (2011). Motivating voter turnout by invoking the self. *Proceedings of the National Academy of Sciences*, *108*, 12653-12656.
- Dimock, M., Keeter, S., Schulman, M., & Miller, C. (2001, May). *Screening for likely voters in pre-election surveys*. Paper presented at the annual meeting of the American Association for Public Opinion Research, Montreal, Quebec, Canada.
- Fast, N. J., Gruenfeld, D. H., Sivanathan, N., & Galinsky, A. D. (2009). Illusory control: A generative force behind power's far-reaching effects. *Psychological Science*, *20*, 502-508.
- Freedman, P., & Goldstein, K. (1996). Building a probable electorate from preelection polls: A two-stage approach. *Public Opinion Quarterly*, *60*, 574-587.
- Gerber, A. S., Huber, G., & Washington, E. (2010). Party affiliation, partisanship, and political beliefs: A field experiment. *American Political Science Review*, *104*, 720-744.
- Gerber, A. S., & Rogers, T. (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *The Journal of Politics*, *71*, 178-191.
- Green, D. P., & Gerber, A. S. (2006). Can registration-based sampling improve the accuracy of midterm election forecasts? *Public Opinion Quarterly*, *70*, 197-223.
- Green, D. P., & Gerber, A. S. (2008). *Get out the vote!* Washington, DC: Brookings Institution Press.
- Highton, B. (2004). Voter registration and turnout in the United States. *Perspectives on Politics*, *2*, 507-515.
- Hillygus, S. D. (2005). Campaign effects and the dynamics of turnout intention in Election 2000. *The Journal of Politics*, *67*, 50-68.
- Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, *39*, 17-31.
- Knack, S. (2001). Election day registration: The second wave. *American Politics Research*, *29*, 65-78.
- Mann, C. (2003, May). *Actions speak louder than words: Are vote likelihood screens effective?* Paper presented at the 2003 Annual Conference of the American Association for Public Opinion Research, Nashville, TN.
- Murray, G., Riley, C., & Scime, A. (2009). Pre-election polling: Identifying likely voters using iterative expert data mining. *Public Opinion Quarterly*, *73*, 159-171.
- Mutz, D. C., & Reeves, B. (2005). The new videomalaise: Effects of televised incivility on political trust. *American Political Science Review*, *99*, 1-16.
- National '12 General Election Poll—Tuesday November 6, 2012. (2012, November 6). *Politico.com*. Retrieved March 1, 2013 from <http://www.politico.com/p/2012-election/polls/president/national/national-12-president-general-election-216>

- Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, *124*, 54-74.
- Payne, J. W., Bettman, J. R., Schkade, D. A., Schwarz, N., & Gregory, R. (2000). Measuring constructed preferences: Towards a building code. In B. Fischhoff & C. F. Manski (Eds.), *Elicitation of preferences* (pp. 243-275). The Netherlands: Springer Netherlands.
- Perry, P. (1960). Election survey procedures of the Gallup poll. *Public Opinion Quarterly*, *24*, 531-542.
- Perry, P. (1979). Certain problems in election survey methodology. *Public Opinion Quarterly*, *43*, 312-325.
- Pinkleton, B. E., Austin, E. W., & Fortman, K. J. (1998). Relationships of media use and political disaffection to political efficacy and voting behavior. *Journal of Broadcasting & Electronic Media*, *42*, 34-49.
- Rosenstone, S., & Hansen, J. M. (1993). *Mobilization, participation and democracy in America*. New York, NY: Macmillan.
- Schacter, D. L. (1999). The seven sins of memory: Insights from psychology and cognitive neuroscience. *American Psychologist*, *54*, 182-203.
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, *8*, 657-661.
- Silver, B. D., Abramson, P. R., & Anderson, B. A. (1986a). The presence of others and overreporting of voting in American national elections. *Public Opinion Quarterly*, *50*, 228-239.
- Silver, B. D., Abramson, P. R., & Anderson, B. A. (1986b). Who overreports voting? *American Political Science Review*, *80*, 613-624.
- Tourangeau, R., Rips, L. J., & Rasinski, K. A. (2000). *The psychology of survey response*. Cambridge, UK: Cambridge University Press.
- Triandis, H. C. (1977). *Interpersonal behavior*. Monterey, CA: Brooks/Cole.
- Wolfinger, R. E., & Rosenstone, S. J. (1980). *Who votes?* (Vol. 22). New Haven, CT: Yale University Press.

Author Biographies

Todd Rogers is an assistant professor of public policy at the Harvard Kennedy School. He is a behavioral scientist whose work bridges the gap between intention and action.

Masahiko Aida is Lead Survey Research Scientist at Civis Analytics. Prior to his work at Civis, he was vice president and chief data scientist at Greenberg Quinlan Rosner Research.