

**Enough Already with Random Digit Dialing:
Can Registration-Based Sampling Improve the Accuracy of
Election Forecasts?**

Donald P. Green
Yale University

Alan S. Gerber
Yale University

November 15, 2003

Earlier versions of this report were presented at the 58th meeting of the American Association for Public Opinion Research, Nashville, Tennessee and the Gallup Conference on Improving the Accuracy of Polling, May 2-4, 2002, Washington, DC. The authors are deeply indebted to Christopher Mann, who played a key role in preparing and analyzing the RBS samples used in this study. We are grateful to Rich Morin and Claudia Deane at the Washington Post, Kathleen Frankovic and Anthony Salvanto at CBS, and Douglas Schwartz at Quinnipiac University, who conducted the polls reported in this paper. We profited from the suggestions of many readers, among them Dale Kulp, Jim Lepkowski, Mark Mellman, and Nick Panagakis. This research was supported by a grant from the Smith Richardson Foundation, which bears no responsibility for the views we express.

Abstract

Prior to the 2002 election, pre-election polls relied almost exclusively on random digit dialing (RDD) to draw representative samples of likely voters. We propose an alternative sampling methodology, stratified random sampling from voter registration lists. Because these lists often furnish useful background information that strongly predicts whether a person will vote, registration-based sampling (RBS) may provide more accurate election forecasts. Sampling from registration lists also reduces the costs associated with identifying likely voters.

In order to test the relative accuracy of RDD and RBS in predicting election outcomes, we collaborated with the Washington Post, CBS News, and Quinnipiac polls, which conducted parallel RDD and RBS surveys in Maryland, New York, Pennsylvania, and South Dakota prior to the November 5, 2002 elections. The results suggest that RBS compares favorably to RDD in terms of forecasting accuracy but at substantially lower cost.

Using Registration-Based Sampling to Improve Pre-Election Polling

Pre-election polling is a highly standardized activity. The overwhelming majority of pre-election surveys rely on telephone interviewing, and the overwhelming majority of phone surveys use random digit dialing, whereby calls are placed to randomly generated phone numbers within pre-selected phone exchanges. Because phone calls are placed to people whose identities and addresses are unknown, the survey analyst must rely on the respondent to furnish all of the key pieces of information needed to interpret the survey results. The respondent must help the interviewer select a randomly chosen member of the household for questioning, who in turn must furnish information about whether they are currently registered to vote and whether they intend to vote in the upcoming election.

Because the aim of pre-election polling is to describe the opinions of the voting public, many respondents are excused during the interview process. Unregistered voters receive a polite thank you, and their interviews are terminated early on. Respondents who make it past this screen are cross-examined about their proclivity to vote in the upcoming election. Those who express uncertainty about whether they will cast a ballot or who report having skipped elections in the past are discreetly removed from the sample afterwards. Remaining respondents constitute the sample of “likely voters” whose opinions (perhaps after a bit of massaging using weights for the respondents’ party affiliation or demographic profile) become the poll’s official results.

Whether this procedure provides a reliable means for forecasting election outcomes depends in part on whether the interviewer accurately discerns which respondents will vote. In many, if not most elections, voters and nonvoters have different candidate preferences; mistaking a nonvoter for a voter in these instances will bias the survey results. Pollsters are, of course, well aware of this problem. When pre-election polls mispredict election outcomes, pollsters and poll-watchers invariably lay the blame at the feet of voter turnout. This ritual is in some ways convenient for pollsters because it means that they need not adjust their survey procedures in order to account for it. The vicissitudes of voter turnout are regarded as a sort of uncontrollable force of nature for which pre-election forecasters are not responsible.

Polling organizations might do a better job of anticipating voter turnout by changing one facet of their standard operating procedure. We propose replacing (or at least supplementing) random digit dialing (RDD) with registration-based sampling (RBS), which is neither difficult nor costly to implement. Indeed, the advent of database management firms, the digitization of records by local registrars, and recent legislation designed to standardize record-keeping on a statewide basis make RBS an increasingly attractive alternative to RDD.

A strong case may be made in the abstract for the advantages of using RBS – and more generally, for the advantages of using background information about respondents for forecasting purposes. The empirical question remains: Does registration-based sampling in fact lead to improved election forecasts? To demonstrate that this alternative

sampling methodology warrants further investigation, we present results from head-to-head comparisons of RDD and RBS in four states.¹ Surveys conducted by the Washington Post, CBS News, and the Quinnipiac Poll during the weeks leading up to the November 2002 elections suggest that RBS is a promising, cost-effective alternative to RDD. RBS performed well in terms of forecasting accuracy, and its accuracy will likely improve as survey researchers learn to take full advantage of the rich contextual information that registration data provide.

The Allure of RDD

During the 1970s, the polling industry began moving away from face-to-face interviewing and embraced phone surveys. The advantages of phone surveys were many, such as lower unit costs, easier supervision, and shorter lead times. At that time, longstanding skepticism about the representativeness of telephone surveys began to subside. The overwhelming majority of American residences had phone service. There remained, however, the problem of drawing a random sample of the population. RDD sampling is premised on the idea that one cannot obtain a representative sample of households using listed phone numbers. Some people, after all, have unlisted phone numbers, and excluding them from the sample leads to bias.

This argument persuaded the polling world to embrace RDD and put up with various problems specific to it. Although the phone digits are drawn randomly, the

¹ Unfortunately, we never learned the results from polls we conducted in three states in collaboration with Voter News Service, which went out of business immediately after the polls were conducted.

people who answer the phone are not. RDD requires the pollster to wade through large numbers of nonworking or nonresidential numbers in search of residences. Once a residence is identified, an accomplishment that requires some kind of contact with a person or an informative answering machine, the pollster must draw a random sample within the household. This task involves either some form of enumeration (e.g., "How many registered voters are living at this address?") or some near-random selection criteria (e.g., "May I speak to the adult who will be having the next birthday?"). Executing this type of within-household random sampling may run into practical difficulties stemming from the suspicion or puzzlement raised by these odd questions.

Even when within-household sampling proceeds without incident, there remains the task of re-weighting the sample to reflect the sampling probabilities. In simple random sampling, each observation has the same probability of being selected. When RDD sampling is clustered by telephone exchanges, these weights are determined by the putative number of residential numbers within each exchange. Another problem is that some residences have more than one number, which means that the pollster must inquire about the number of phone lines in the household that could be used in a survey conversation (e.g., phone lines not dedicated to a computer modem). The final survey data are then re-weighted according to the probability of selection into the sample.

Were the survey simply attempting to draw a representative sample of the general population, these complications might be worth putting up with in order to reap the main benefit of RDD, namely, overcoming the problem of unlisted residential numbers. RDD

seems to do a reasonable job of sampling the adult population. But when the aim is to draw a representative sample of those who will vote on Election Day, the merits of RDD become more ambiguous.

Sampling the Population of Actual Voters

Ideally, a pre-election poll is a random draw from the population of people who will actually cast ballots on Election Day. To be sure, minds may change after the poll is conducted, but error associated with opinion change is conceptually distinct from the error associated with sampling from a population other than that of actual voters. Even before Election Day, we have quite a bit of information about who will vote. The population of actual voters is a subset of the population of registered voters. With the exception of a few states that have same-day voter registration or lack registration requirements, the population of voters is drawn from the population of those who have registered at least a few weeks before the upcoming election. Incidentally, in states with traditional registration systems, only a tiny fraction of the voting electorate registers during the months leading up to an election. For example, in Raleigh, North Carolina, 7% of the 264,669 voters who cast ballots in 2000 had registered after October 1, 2000. In Seattle, the corresponding figure is 4% of 786,286.

Knowing which voters in an RDD survey are registered gives us a leg up on the problem of describing the population of those who will actually cast ballots. When one uses RDD, one typically relies on the respondent to indicate whether he or she is

registered. One could conceivably perform a reverse phone match linking the RDD number to a registration list, but the phone numbers on registration lists tend to be outdated, many registrars do not require phone numbers, and the respondent must furnish address information in order to find the right registration list (Anderson and Saad 1996; Dimock et al. 2001).

Figuring out who is registered is really just the first step toward knowing who will actually vote. Here, too, we must rely on the representations made by the respondent. The typical pre-election survey peppers respondents with questions about the frequency with which they voted in the past and their likelihood of doing so in the upcoming election. Based on this series of questions, the survey analyst makes some kind of determination about which respondents constitute "likely voters." (Alternatively, the analyst could assign each respondent a probability of voting, but this does not seem to be the industry practice.) Survey results are tallied for this pool of likely voters, and that's that.

The problems with this approach are twofold. First, respondents may misstate their voting intentions. Second, the survey analyst is generally at a loss to translate even a sincere report of past or intended voting behavior into a probability of voting in the upcoming election. Naïve classification of these responses into rigid "likely voter" and "non-likely voter" categories only compounds the problem, since it discards meaningful variation within these groups. Taken together, these problems mean that the "likely voter" designation is really just guesswork.

Interestingly enough, RBS can provide information that may help improve RDD voter screening methodology. The relationship between vote intention and actual voter turnout is easily gauged for RBS subjects.² (So, too, is the relationship between nonresponse and voter turnout.) The researcher draws a sample from a list of registered voters, conducts a phone interview, and examines the relationship between vote intention measures and actual voting behavior. Indeed, by looking after the fact at the pre-election preferences of only those respondents who actually voted, one can gauge the forecasting error of the survey net of uncertainty about voter turnout. The recent work by Mann (2003a, b) illustrates the value of using RBS surveys to gauge the predictive accuracy of voter screens. Mann also shows that past voter history, which is available to RBS, is as valuable as stated vote intention in terms of forecasting whether people actually vote.

Registration Based Surveys

One of the main advantages of sampling from registration lists is that they contain a wealth of information that may be used to forecast voter turnout. Although the quality of information available from public records varies across jurisdiction, the typical registration file contains date of birth, date of registration, party registration, and the

² Scholars frequently look to the American National Election Study surveys for this information, but these surveys have not performed vote validation in recent years, and older surveys tended to rely on in-person interviewing, which may have different misreporting rates than phone surveys. Validated vote studies conducted by commercial firms doing pre-election phone surveys are rarer still, and those conducted to date have run into operational problems when determining voter turnout (Anderson and Saad 1996; Dimock et al. 2001).

number of voters registered at each address. In addition, one may obtain data on past voter turnout. Sometimes these voter histories are furnished by registrars, and sometimes by private vendors. The reliability of these data doubtless varies, but the point remains that one can produce a powerful statistical prediction of future voting behavior simply by reference to public records -- before one even speaks with a respondent.

The strong statistical relationship between public information and voter turnout is illustrated by voting patterns in Raleigh and Seattle. Suppose we were interested in forecasting voter turnout in the 2000 general election. Our voter history files contain two very useful pieces of information: when the voter registered and whether he or she voted in the 1998 general election and the 2000 primary. Table 1 presents the cross-tabulation of voter turnout in 2000 by whether the individual was registered and voted in the preceding 1998 and 2000 elections. Table 1 shows that a voter's prior vote history predicts their subsequent voter turnout. For example, in Seattle we see that among the 411,526 voters who were registered in 1998 but skipped both the 1998 and 2000 primary elections, turnout in the 2000 election was a dismal 31.0%. By contrast, those who voted in both these previous elections (N=326,302) voted at an astonishing rate of 97.4%.

Intuition suggests that any survey of the voting electorate should place special emphasis on the opinions of citizens with a high ex ante probability of voting. Consider the limiting case, for example, in which members of the population come in two types, those who are certain to vote and those who are certain to abstain. There would be no point in interviewing any members of the latter group. By this logic, constructing an

optimal sample – that is, a sample that gives the smallest prediction errors when forecasting the actual election outcome – should take notice of prior voting history because it predicts whether someone will vote in the upcoming election.

Current RDD sampling procedures are a long way from optimal. Pollsters either discard or down-weight the interviews they conduct with those they deem to be unlikely voters. Interviewing nonvoters is a waste of resources, but RDD does not give pollsters much choice, because there is no way to tell which respondents are likely to vote before interviewing them. RBS, on the other hand, gives pollsters a means of differentiating between likely voters and nonvoters *before* money is spent interviewing them.

At this point, it may be helpful to distinguish between two different variants of RBS. Simple random sampling from registration lists, a technique that is common among political pollsters and occasionally by academic researchers (e.g., Visser et al. 2000; Chang and Krosnick 2001), might be dubbed random RBS, or RRBS. This approach is not optimal, because it makes no use of information such as past voting history that can be very helpful in predicting whether a potential respondent will vote in the upcoming election. When we think of RBS, we have in mind something rather different. Each person on the registration list is assigned a probability of voting in the next election, based on factors such as registration date and past voting history. These probabilities enable the researcher to create a profile of those who will actually cast ballots. A random sample is drawn from this expected electorate.

In practice, the procedure boils down to the following steps. First, one breaks the registration list into strata. For purposes of illustration, let us divide our registration list into three strata: those who were not registered prior to the last election, those who were registered but did not vote, and those who voted. Second, one estimates the rate at which each stratum is expected to vote (more about how to do this in a moment). Third, one multiplies the voting rates times the number of people in each stratum in order to find the size of the anticipated electorate and what fraction of the electorate each stratum constitutes. These fractions then become the sampling weights for the RBS survey. Imagine that the three strata consist of 100,000 new registrants, 400,000 past abstainers, and 500,000 past voters. Suppose that voting rates among the three strata are 50%, 20%, and 90%. The expected electorate therefore consists of 50,000 new registrants (8.6%), 80,000 past abstainers (13.8%), and 450,000 past voters (77.6%). When drawing a sample from the registration list to be interviewed, 8.6% of the sample should be new registrants, 13.8% past abstainers, and 77.6% past voters (see Thompson 1992, Chapter 11).

Of course, in practice, voting rates within each stratum are not known beforehand. It is therefore handy to know how voting history has predicted subsequent turnout in recent elections. Before attempting to construct an optimal survey design for the 2002 election, we examined the relationship between turnout in the 1996 and 1998 elections. We found, for example, that Pennsylvanians who voted in the 1994 and 1996 elections comprised 60% of the 1998 electorate. In the RBS survey for 2002, therefore, 60% of the names selected for our sample consisted of voters who had voted in the two previous

federal elections. Luckily, actual voter turnout figures for 2002 show that 61% of those who cast ballots had voted in the two previous federal elections. We were less fortunate in the case of New York, where we anticipated that 48% of the electorate would be comprised of those who voted in the two previous federal elections. It turned out that 62% of the electorate consisted of these regular voters. As we will see below, our New York sample under-represented stalwart voters, who took a dim view of independent candidates.

As pollsters develop more extensive data bases and models for forecasting turnout based on past voting records, these errors will diminish. For the present, it should be stressed that getting the rates just right is seldom crucial. Knowing simply that one group is much more likely to go to the polls than another helps a pollster allocate resources more efficiently than would be the case with simple random sampling. Even if the eventual sample is not optimal, it is still a big step in the right direction.

Like any sampling scheme, stratified RBS can be made more complex. One could, for example, weight observations according to background characteristics such as age, party registration, gender, or the political characteristics of their voting precincts.³

³ A more sophisticated weighting scheme would take account of the fact that the candidate preferences of voters with ardent partisan preferences are easier to predict than those with moderate views. Knowing that a voter is a regular participant in Republican primaries leaves relatively little uncertainty about how he or she will vote. For this reason, the pollster need not interview as many regular primary voters as voters in other equally sized centrist groups in the population. The statistical rule of thumb is this: allocate more interviews to large subgroups of voters, unless you know ahead of time how they are likely to vote. For example, in Seattle, primary voters who also voted in 1998 comprised 41.5% of the 2000 electorate. This group requires a bit of attention due

The more interesting question is what to do with stated vote intention. The issue is whether screening RBS surveys for self-described “likely voters” or imposing other weighting schemes based on stated vote intention improves the forecasting accuracy of RBS polls. Although we provide a tentative answer below, we hasten to add that the technology of RBS forecasting remains in its infancy. As experience with RBS forecasting accumulates, pollsters will learn to develop more effective ways of melding registration-based data with survey responses.

Although we have used the term “registration” to refer to voter registration lists, the idea behind RBS extends to any sort of list-based sampling frame. For example, in polities without voter registration, postal lists could be used as the basis for sampling. The statistical advantages of RBS, however, stem from the ancillary information about the prospective respondent that one gathers before interviews are conducted. Voter registration lists at a minimum tell us the voter’s age, date of registration, and whether other voters are registered at the same address – all important predictors of voter turnout. Moreover, registration lists include address, which enables the pollster to find out a great deal about the political and economic climate in which the respondent lives and, by extension, how the potential respondent is likely to behave on Election Day.

to its enormous size. On the other hand, a pollster will quickly learn after a few interviews that the GOP stalwarts favor Bush and that the Democratic regulars favor Gore. One could invest additional interviews in this group, but those interviews are better spent on other large groups about which there is more uncertainty. The next largest group, those who voted in 1998 but skipped the 2000 primary, comprised 26.6% of the 2000 electorate. This group would command extra polling resources.

RBS in Practice

Having made the case for registration-based sampling, we now turn our attention to the practical aspects of sampling from registration lists and eliciting cooperation from respondents.

Obtaining registration lists. Registration-based sampling starts with a registration list. These lists are, with very few exceptions, available from one or two sources. The first source is the registrar of a local jurisdiction, such as a city or county. Although we cannot claim to have done an exhaustive study of the extent to which these local registrars work with digitized lists, we have experience with local registrars. Our voter mobilization experiments have made us pen pals with registrars across the country -- in small rural communities, suburbs, and large cities. We have rarely encountered registration lists that are not in machine-readable form. Typically, all that is involved in obtaining this list is a small fee, although sometimes one must also promise not to use the list for telemarketing.

In some cases, these lists are officially off-limits to anyone who is not connected with a political party. Parties seem to be quite porous, however, and these forbidden lists are commercially available from database vendors. These vendors charge more for these data than local registrars, but the prices are still quite modest. The advantage of purchasing these data from vendors is that they have often taken the trouble to append to the dataset information about the previous elections in which a person has voted;

registration lists obtained directly from registrars may not contain this information, although it can often be obtained from them separately in machine-readable form. The quality of these lists varies across states and sometimes within them. Until the federally mandated standardization of voter registration lists takes effect in 2006, pollsters who wish to use RBS must investigate the feasibility of this approach for the regions they seek to study.

Fortunately, pollsters are typically most interested in close races, which are precisely the races that most interest political parties and campaigns. Parties and campaigns are the list vendors' big customers, and vendors make special efforts to keep up-to-date registration lists for tight races. Unfortunately, vendors, like pollsters, are sometimes caught off-guard by a race that suddenly heats up. As we note below, this type of surprise occurred in Maryland, when the governor's race suddenly became competitive. As a result, the registration list that we used did not have vote history for the 2000 election. We were forced, therefore, to stratify based on 1994 and 1996 vote history rather than on the 1998 and 2000 vote history. Not optimal, but still serviceable.

Matching addresses and phone numbers is the point at which slippage occurs in the RBS sampling process. The availability of listed numbers varies widely by region. Iowa is very good; California is very bad. In some cases the matching rate can be improved by sending registration lists out to multiple vendors, who draw their phone numbers from different sources. The list vendor that supplied the statewide registration lists for Maryland, New York, Pennsylvania, and South Dakota also supplied valid phone

numbers for approximately two-thirds of each sample.⁴ Nevertheless, failure to obtain numbers from one-third of the electorate is an important concern, particularly if unlisted voters have distinctive political views. The empirical test presented below is motivated primarily by our desire to gauge the severity of this bias.

Of course nothing prevents the pollster from falling back on RDD in cases where the phone match is particularly poor. It should be stressed, however, that the forecasting accuracy of RBS in comparison to RDD is an empirical question, for as we note below, RBS may have certain advantages in terms of reducing nonresponse. If RDD phone surveys are enlisting the cooperation of less than half of the households they contact, it is by no means clear *ex ante* which source of bias is more problematic. The only way to assess the relative accuracy of RDD and RBS is to do both and see what happens. As experience with RBS accumulates, pollsters may develop hybrid approaches that use weighted averages of both RBS and RDD results.

Telephone interview protocol. Once the RBS sample of respondents has been selected, the mechanics of the survey are quite straightforward. Since the identity of the respondent is known to the interviewer, the caller may ask for the respondent by name. No further enumeration of the household is necessary.

A useful illustration of how response rates may vary by sampling methodology comes from a Washington Post survey of the Maryland governor's race conducted

⁴ The rates were 65% for Maryland, 69% for New York, 66% for Pennsylvania, and 70% for South Dakota.

approximately 9 days before the 2002 general election. RDD and RBS polls were conducted by a single calling house over an identical time period. The surveys were the same, aside from the fact that the RBS poll asked for respondents by name and dispensed with the within-household respondent-selection procedures used in the RDD survey. The RDD poll attempted to reach residents at 14,051 distinct phone numbers and completed interviews with 1,738 persons (12%). Among the 960 who claimed to be registered to vote, 725 claimed to be “certain” to vote in the upcoming election. The RBS poll attempted to reach 3,754 registered voters and completed interviews with 838 (22%), of whom 657 claimed to be certain to vote.⁵

While knowing the names of potential respondents alleviates some of the clumsiness of RDD enumeration procedures, it may have some drawbacks as well. It appears that respondents to RBS surveys may need extra assurances about confidentiality. One subtle but nonetheless noteworthy difference between RBS and RDD response patterns is that the former show more reluctance to disclose vote intentions. Across the eight elections the average percentage who declined to express a candidate preference is slightly higher in RBS samples (18.9%) than in RDD likely voter samples (17.3%). However, Mann (2003a) presents evidence suggesting that reluctance to respond depends on whether respondents received an advance letter. Those who received a letter were slightly *more* likely to voice a candidate preference than RDD likely voters, although the difference is not statistically significant.

⁵ In terms of gaining the cooperation of respondents, one advantage of RBS is that it enables survey firms to send out a letter ahead of time explaining the purpose of the survey, inviting participation, and perhaps offering some incentive. Mann’s (2003a) analysis of RBS surveys in Maryland, New York, and Pennsylvania indicates that sending a letter to respondents ahead of time significantly increases cooperation and completion rates.

The Performance and Cost Effectiveness of RBS Polls

In order to compare the forecasting accuracy and cost of RBS and RDD surveys, we collaborated with three polling organizations – CBS News, Washington Post, and Quinnipiac University – which conducted statewide polls in South Dakota, Maryland, Pennsylvania, New York.

CBS News conducted the South Dakota poll from October 9-11. The RDD and RBS surveys were both conducted using a CATI system to automate calling and coding of responses. Up to nine attempts were made to contact respondents.

From October 20-24, the Washington post conducted simultaneous surveys using RDD and RBS samples in Maryland. The surveys were conducted by TNS Intersearch, the firm which normally conducts polling for the Washington Post. Both surveys were conducted simultaneously using a CATI system to automate calling and coding of responses. Up to eight attempts were made to contact respondents.

The Quinnipiac University Polling Institute (see Schwartz and Richards 2003) conducted simultaneous surveys using RDD and RBS samples in Pennsylvania during October 21 –27. From October 28 to November 3, the Institute also conducted RDD and RBS surveys in New York State. The pairs of surveys in each state were conducted using the facilities and staff at Quinnipiac University. The RDD surveys were conducted using

a computerized CATI system to automate the process of calling and completing the questionnaire for live interviewers. Due to capacity limits, the RBS surveys were conducted using paper questionnaires filled out by the interviewers, who also manually dialed the phone numbers. On weeknights, calls were attempted from 5:30pm to 9:30pm for the RDD survey and from 7:00pm to 9:30pm for the RBS survey. On weekends, calls were attempted during the same periods for the RDD and RBS surveys: 10:00am to 3:00pm on Saturdays and 4:00pm to 9:00pm on Sundays. The management of the sample for the paper RBS survey was kept as similar as possible to the computerized RDD sample.

The RDD and RBS surveys in each state used nearly identical questionnaires. In each state the surveys differed in their introduction: the RBS questionnaire requested the voter by name, while the RDD questionnaire asked for the individual over 18 with the next birthday (a pseudo-random selection technique frequently used in RDD polls). In New York, the survey asked voters whom they would vote for in the race for governor, attorney general and comptroller. South Dakota's survey gauged voter preferences for Governor, Senate, and House. The Pennsylvania and Maryland surveys covered only the Governor's race in each state. The vote preference question asked whether voters were certain to support the candidate or might change their mind. If voters declared no preference in the initial question, they were asked if they leaned towards a candidate; the results reported below include the preferences of leaners. In Maryland, the questionnaires had another major difference: the RDD questionnaire followed the vote preference section with questions about the favorability ratings of the candidates and

about a number of issues in the campaign before concluding with a battery of demographic questions common to both questionnaires; the RBS questionnaire omitted the favorability and issue questions. In New York and Pennsylvania, both the RBS and RDD questionnaires followed the vote preference questions with questions on the favorability ratings of the candidates and concluded with a battery of demographic questions. Since we focus our attention on vote preference, the fact that the questionnaires diverge after the vote preferences questions were asked is inconsequential.

The RDD samples in each state were provided by Survey Sampling Inc. [SSI], who generated numbers randomly for the area codes within the state. SSI attempted to purge random numbers that were not in service from the potential sample. The RBS samples in each state were drawn randomly from a list of registered voters in each state maintained by Voter Contact Services (VCS). During the late summer, VCS gathered current lists of registered voters from each county of the three states. VCS completed the updating their list of registered voters in New York and Pennsylvania by August 21st, 2002. Because a few counties (4.2% of the registered voters) did not collect voter history in the early 1990s, we used a simplified stratification system that provided only 3 strata rather than our usual 6 in Pennsylvania. Due to delays in re-districting and the late primary election, VCS was not able to gather the 2000 voter history prior to the election in Maryland, necessitating an adjustment in our stratification methodology to use elections 4 and 6 years prior rather than 2 and 4 years prior as we had done in New York and Pennsylvania. In Maryland, the registration list was updated on September 5, 2002. The voter records in Maryland, Pennsylvania, and New York were matched to phone

records that were current as of October 2002. South Dakota registration data were updated as of September 23 and obtained directly from the state. The registration information was matched to a database of phone numbers by InfoUSA.

Prior to drawing the RBS sample, the entire list of registered voters was stratified into mutually exclusive groups based on past voting history. After estimating what share of the anticipated electorate each stratum represented, we drew a stratified sample of 40,000 registered voters in each state. From this stratified random sample, a sub-sample of 10,000 voters with phone numbers were selected in each state. Half of the RBS samples in Maryland, New York, and Pennsylvania were send mail in advance explaining the purpose of the survey and encouraging their participation. As Mann (2003a) reports, the letter increased survey response rates but did not lead to a consistent improvement in forecasting accuracy. We therefore pool all of the RBS responses for a given state for purposes of the analysis presented below.

RBS vs. RDD: Predictive Accuracy in Forecasting 2002 Election Outcomes

The central question of this study is whether the use of RBS increases forecasting accuracy. It could be argued that the reliance on listed telephone numbers introduces bias. Indeed, the Maryland RBS poll shows indications of under-representation of African Americans (Deane and Morin 2003), while the Pennsylvania and New York RBS polls show some under-representation of urban dwellers (Schwartz and Richards 2003). On the other hand, it could be that RBS outperforms RDD. RBS generates more usable

interviews, and its weighted sampling frame takes advantage of hard data on respondents' prior voting behavior. The magnitude of these biases therefore remains an empirical question.

Tables 2 and 3 present the head-to-head comparison of RBS and RDD for the eight races about which vote intention was measured in our four statewide surveys. The leftmost columns of Tables 2 and 3 list the type of office as well as candidates' names, party, and incumbency. Each of the four state polls involves a gubernatorial race. The South Dakota poll also involved reasonably close races for House and Senate. The New York poll solicited opinions about a lopsided attorney general race and a more competitive, but obscure, comptroller election. The actual vote outcome for each race is presented for all candidates receiving more than 5% of the vote.

The next four columns of Tables 2 and 3 present the RBS results under different weighting and sampling criteria. The first of these columns presents the unweighted results using all RBS respondents. Note that "unweighted" in this context means that no weighting was done after the samples were created and given to the survey organizations. Recall that these stratified RBS samples were initially created using weights for voting propensities and were designed, therefore, to be representative of the voting population. So "unweighted" in this context means that no further weighting was used after the surveys were conducted.

The second column weights the RBS data to take into account the fact that different sampling strata had different survey response rates. Those who completed interviews were weighted so that the initial strata proportions were preserved. The third column restricts the RBS sample to “likely voters,” that is, respondents who (depending on the wording of the question) claimed that they would “definitely” vote or were “certain” to vote.

Finally, we restricted the RBS sample to those respondents who actually voted, according to voter turnout records. Obviously, isolating this subgroup is of no practical value to pollsters seeking to anticipate election outcomes, but it serves an important methodological function. By comparing results for actual voters to other RBS weighting approaches, we can assess the value-added of an optimal turnout forecast. In addition, by comparing actual voters in the survey to actual voting outcomes, we can gauge each survey’s sampling error without any turnout-related error.

For RDD surveys, two sets of results are presented. The first column presents the raw survey results tabulated for all RDD respondents. The second column restricts attention to “likely voters” as defined by each polling organization. In Maryland and South Dakota, likely voters were determined based on a single question; in Pennsylvania and New York, an index based on several questions was used instead. The results for likely voters are also weighted by various demographic factors such as age and gender. For Maryland, New York, and Pennsylvania, these numbers are the final results reported for public release by each organization. Because the South Dakota RDD poll was not

reported as part of a news release, we calculated these results based on the raw data for likely voters.

The predictive accuracy of these polling results may be gauged using two different approaches. The first uses the vote margin,⁶ which is calculated by subtracting the percentage preferring the Republican from the percentage preferring the Democrat, as in Table 2. The second approach, used in Table 3, calculates the average forecast error for each candidate, after allocating undecided voters. Undecided respondents were assumed to split their votes in the same way as those expressing a candidate preference. Average forecast error is calculated by taking the absolute value of the difference between each candidate's forecasted percentage and the actual percentage that this candidate received in the election.

Table 2 reveals that unweighted RBS provides surprisingly accurate forecasts of the actual vote margin. In 6 of the 8 races, unweighted RBS outperforms RDD polls that use a likely voter screen. One race in which RDD outperforms RBS is the lopsided 3-way race for New York Governor, and here the principal advantage of RDD is its superior prediction of the vote won by the third-place finisher. In the remaining races, unweighted RBS performs well.⁷ The other race is the South Dakota gubernatorial

⁶ Note that the standard error of the vote margin is twice the standard error of a single candidate's vote share. For this reason, the National Council on Public Polls recommend dividing the forecast error of the vote margin by two. For the sake of simplicity, we have not made this adjustment here.

⁷ As Morin (2003) points out, tinkering with the unweighted Maryland sample to adjust for its supposed under-representation of African-American respondents does not improve

contest. Here, RBS overestimates the Republican lead by 3.0 points, while RDD underestimates it by 2.7 points. Overall, the average absolute error is 3.1 percentage-points for RBS and 6.4 percentage-points for RDD when predicting actual vote margins.

Table 3, which allocates undecided respondents when candidate vote percentages, again suggests the superiority of RBS. Overall, the average absolute candidate error for unweighted RBS (2.8) is lower than for RDD samples of likely voters (3.9). Excluding the two races for lower offices in New York brings these figures down to 1.4 for RBS and 2.4 for RDD. Interestingly, O’Neill, Mitofsky, and Taylor (2002) report that the average absolute forecast error for Gubernatorial and Senatorial candidates of 159 polls conducted after October 19, 2002 was 2.4 percentage-points. Thus, the RDD polls in our study performed as well as RDD polls generally do, but RBS polls did somewhat better.

The forecasting edge of RBS in the four states we studied stems in part from the fact that each of RDD samples of likely voters has approximately 100 fewer respondents. The contrasting sample sizes reflect the greater efficiency of RBS. The RDD samples of registered voters are larger than the corresponding RBS samples, but the likely voter screen causes this number to drop sharply. RDD surveys that use the likely voter screen have smaller forecasting errors than RDD surveys that screen only for registered voters, but RBS outperforms them both.

the poll’s predictive accuracy. The lesson from the array of forecasting results presented here seems to be that one should leave well enough alone.

Interestingly enough, the forecasting accuracy of RBS does not improve noticeably when we correct the RBS strata for response bias or focus attention exclusively on respondents who claim to be likely voters. As Mann (2003b) points out, patterns of survey nonresponse are in some sense their own likely voter screen; those who respond to surveys are disproportionately drawn from the ranks of regular voters. The strong performance of unweighted RBS is fortunate for survey researchers, because it means that they may dispense with cumbersome weighting procedures or the costly practice of discarding “unlikely” voters.

Nor does forecasting accuracy improve when we limit the sample to those RBS respondents who actually voted. RBS nonvoters are sufficiently similar to RBS voters that the loss of those observations reduces bias only minimally, while increasing sampling error.

Needless to say, one cannot form a definitive judgment about the superiority of RBS based on just 4 polls covering 8 elections. We do not find any indication, however, that RBS suffers from debilitating biases as the result of its failure to solicit the opinions of those without listed phone numbers. Nor do we find any systematic tendency for RBS to exaggerate the support for candidates of one party. In its first direct empirical test, RBS acquitted itself reasonably well and clearly warrants further experimentation in the future.

Cost Comparison

When comparing the costs of RDD and RBS polls, we shall assume that the polls are conducted in the same manner (e.g., computer-assisted telephone interviewing), by the same staff, using a questionnaire of similar length. The two polls differ in cost as a function of expensed incurred while creating a sample, recruiting respondents, and interviewing. Currently, RDD enjoys a cost edge in terms of sampling. Obtaining random phone numbers is relatively inexpensive. For example, the Maryland sample of RDD numbers cost \$614, whereas the New York and Pennsylvania samples cost \$1200 apiece. Commercial firms have yet to offer stratified samples from registration lists. The average cost of the registration lists in each state was approximately \$1900 in each state, including the costs of the phone match.⁸ This cost would be somewhat higher if we had not performed the stratification ourselves.

The costs of recruiting respondents depend on whether an RBS survey is preceded by a letter encouraging respondents to participate. This letter can be costly to print and mail. The cost of sending the letter to roughly half of the RBS target list in each state was \$1956 in Pennsylvania, \$2003 in New York, and \$2500 in Maryland.⁹ As we discovered during our ill-starred attempt to poll the New York gubernatorial primary, the

⁸ This figure does not count the extra names that we purchased but consigned to an uncontacted control group. The control group was to be used as part of a separate study designed to assess whether those who are surveyed are more likely to vote as a consequence. The cost breakdown for each state was \$1489 (Pennsylvania), \$1773 (New York), \$1474 (Maryland), and \$2871 (South Dakota).

⁹ The Maryland letter was sent first-class, because the survey followed soon after the mailing list became available.

expense of the letter may go for naught if a leading candidate withdraws from the election. Advance letters were used for half of the target samples in Maryland, New York, and Pennsylvania.

If one foregoes the advance letter, the balance sheet clearly favors RBS, because an RDD poll devotes a great many hours sifting through nonviable phone numbers in search of respondents. The interview process is where the cost savings of RBS become most apparent, because one need not screen for registered voters or, evidently, likely voters. TNS Intersearch, which conducted the Washington Post Poll in Maryland, estimated that the cost per completed interview with a registered voter averaged \$44.91 for RDD and \$22.19 for RBS (Morin 2003). For sample sizes of 750 respondents, this amounts to a savings of approximately \$17,000. The Quinnipiac Poll measured efficiency in terms of total interviewer hours associated with each poll. RDD (using CATI) consumed 633 hours in Pennsylvania and 510 in New York. RBS (using paper questionnaires) consumed 365 hours in Pennsylvania and 295 in New York. RBS and RDD generated comparable numbers of likely voters, but the labor costs required to do so were 40% lower for RBS.¹⁰

The cost advantages of RBS are magnified when pollsters attempt to study specialized populations. In most states, where absentee voters constitute a small but important fraction of the electorate, it is prohibitively expensive to conduct an RDD survey of absentee voters. If, for example, 50% of the adults in a state vote in federal

¹⁰ Cost efficiency estimates are not available for South Dakota.

midterm elections, and only 20% of all voters cast absentee ballots, RDD must plow through approximately 10 screening interviews in order to reach a single absentee voter. Similar arguments apply to surveys of young voters, minority voters, or voters who have recently switched parties. RBS seems particularly valuable for journalists or political marketers intent upon studying small segments of the electorate.

A similar point holds for populations defined by geographic boundaries that do not coincide with telephone exchanges. Consider, for example, the case of the 5th Congressional district in Connecticut. In the wake of redistricting, the borders of this district changed substantially, and one could not necessarily count on RDD respondents to report accurately whether they resided in the newly drawn district. Even if respondents could be trusted to reliably identify their Congressional district, the RDD poll would have incurred additional costs when screening ineligible respondents. Address information from the registration rolls, on the other hand, could be fed into GIS software containing the boundary files of the new district. Thus, the Quinnipiac Poll was able to gauge opinion in the 5th district during the 2002 campaign, generating an accurate and affordable pre-election forecast.

The cost savings from RBS should properly be seen as contributing to the accuracy of RBS polls. Resources that would otherwise be spent on discarded interviews with unlikely voters are freed up for other purposes. A mailing to potential respondents is one option; larger sample size is another. The apparent forecasting advantages of RBS

presented above are in some sense understated, because they do not take into account the ways in which reallocated resources might enhance RBS surveys.

Conclusion

Random digit dialing shares much in common with the QWERTY keyboard. Both inventions grappled with an important practical limitation at the time they were conceived. The QWERTY keyboard resolved the problem of typewriter keys sticking together when the typist went too quickly. RDD overcame the problems of unlisted residential phone numbers and spotty machine-readable lists of eligible voters. Both QWERTY and RDD have become nearly universal, but neither has adapted to other technological changes. Computer keyboards can accommodate much faster typing speeds, and digitized registration lists abound. There is no end in sight for the QWERTY keyboard because so many people are used to working with a certain layout of keys. The investment in RDD is different. Polling firms make extensive use of RDD, but they also conduct surveys of specific populations, such as consumers or business executives. There is more room for change, and some polling firms might welcome the opportunity to get out from under the inconvenience of phoning numbers at random.

This point has special resonance as pollsters look ahead to the party primaries of 2004. The cost advantages of RBS polling in closed party primaries are overwhelming. Costly as it is to locate voters in federal midterm elections using RDD, it is far more difficult to locate Democratic primary voters. Registration databases typically provide

information about who has voted in past primaries – a strong predictor of who will turnout to vote in the upcoming primaries.

The next challenge for registration-based sampling is surveying the national electorate. Faced with registration lists of varying quality across the country, the survey researcher naturally turns to clustered probability sampling. Here the standard logic of clustered sampling applies. Create primary sampling units, and choose randomly among them. Then break the primary units into secondary sampling units, and so forth. Eventually, one gets down to small clusters, such as a Census block group. If, for example, one were to select 200 clusters with 50 names apiece, one would have a working sample of 10,000 names, which would be ample to support a phone survey of 1,000 completed interviews even assuming a mediocre response rate. While gathering registration data from these clusters is a nuisance (although a nuisance that database vendors would happily embrace, for a fee), creating a bit of sampling infrastructure has its benefits. One can sample in the future from clusters located in these towns, merely updating the existing database with new registrants and new voter histories. One could even contemplate panel studies or analyses that made use of contextual information provided by other respondents living in the same local cluster.

The prospects of a national survey with contextual data invite survey researchers to rethink the way that pre-election polls are conceived and analyzed. One interesting possibility is that RBS polls will eventually be treated in a manner analogous to exit polls. Analysts of exit polls forecast the overall distribution of the vote by using voting

returns from selected precincts whose historical voting patterns have been calibrated to past election outcomes. Knowing that a historically Democratic precinct has given lukewarm support to a Democratic candidate means that the Republican opponent is very likely to win. In much the same vein, RBS surveys may be thought of as a collection of small random samples of precincts, whose historical voting patterns are known. Knowing that Democrats from historically Democratic precincts are expressing lukewarm support for the Democratic candidate bolsters one's expectation of a Republican victory. Again, it remains to be seen whether contextual political information improves the forecasting accuracy of RBS surveys, but it is alluring to think that pollsters have untapped opportunities for improving their predictions.

Although this paper has proposed some radical revisions in the way that pre-election polls are designed and analyzed, it should be stressed that we are not calling for an end to RDD, even within the domain of pre-election polling. Our thesis is simply that registration-based sampling warrants further study, refinement, and rigorous evaluation. The results presented here suggest that RBS has the potential to improve accuracy while reducing costs. Time will tell whether RBS is indeed a better mousetrap.

References

Anderson, Andy, and Lydia K. Saad. 1996. Checking up on Respondents: A Voter Validation Study of the 1992 General Election. Paper presented to the 51st Annual meeting of the American Association for Public Opinion Research, Salt Lake City, Utah.

Chang, LinChiat, and Jon A. Krosnick. 2001. Improving Election Forecasting. Unpublished manuscript, Ohio State University.

Deane, Claudia, and Richard Morin. 2003. Polls in Black and White: Examining the Differences in the Demographics of RDD and RBS Surveys in Maryland. Paper presented to the Annual Meeting of the American Political Science Association, Philadelphia, Pennsylvania.

Dimock, Michael, Scott Keeter, Mark Schulman, and Carolyn Miller. 2001. Screening for Likely Voters in Pre-Election Surveys. Paper presented to the 56th Annual meeting of the American Association for Public Opinion Research, Montreal.

Mann, Christopher. 2003a. Getting Pre-Election Surveys Right: The Effects of Advance Letters on Pre-Election Forecasting. Paper presented to the 58th Annual meeting of the American Association for Public Opinion Research, Nashville, Tennessee.

Mann, Christopher. 2003b. Improving Pre-Election Forecasts from Registration Based Sampling: Using Voter Registration Data to Predict Partisan Vote Intention and to Allocate Undecided Voters. Paper presented to the Annual Meeting of the American Political Science Association, Philadelphia, Pennsylvania.

Morin, Richard. 2003. Smackdown in Maryland: RBS versus RDD. *Public Perspective* 6:(January/February) 7-9, 41.

O'Neill, Harry, Warren Mitofsky, and Humphrey Taylor. 2002. National Council on Public Polls Polling Review Board Analysis of the 2002 Election Polls. Report posted at www.nccp.org, December 19, 2002.

Schwartz, Doug, and Clay F. Richards. 2003. The Big East Two: A Comparison of RBS and RDD Polls in the 2002 Elections in New York and Pennsylvania. Paper presented to the Annual Meeting of the American Political Science Association, Philadelphia, Pennsylvania.

Thompson, Steven K. 1992. *Sampling*. New York: John Wiley & Sons.

Visser, Penny S., Jon A. Krosnick, Jesse F. Marquette, and Michael F. Curtin. 2000. Improving election forecasting: Allocation of undecided respondents, identification of likely voters, and response order effects. In Paul Lavrakas & Michael Traugott (Eds.), *Election Polls, The News Media, and Democracy*. New York, NY: Chatham House.

Table 1

Voter Turnout in 2000 by Voter Turnout in Two Previous Election, Raleigh and Seattle

| Raleigh | Vote 2000 | Not Registered in 1998 | | | Abstained in 1998 | | Voted in 1998 | |
|-----------------------------------|--------------|---------------------------------------|------------------------------|--------------------------|------------------------------|--------------------------|------------------------------|--------------------------|
| | | Not Registered for 2000 Primary | Abstained 2000 Primary | Voted 2000 Primary | Abstained 2000 Primary | Voted 2000 Primary | Abstained 2000 Primary | Voted 2000 Primary |
| Voted in 2000 General Election | NO | 12,053 (34.7%) | 44,159 (67.3%) | 236 (5.6 %) | 60,631 (49.8%) | 363 (6.0%) | 12,137 (11.7%) | 1,363 (2.3%) |
| | YES | 22,622 (65.3 %) | 21,483 (32.7 %) | 3,867 (94.4%) | 61,038 (50.2%) | 5,691 (94%) | 91,738 (88.3%) | 59,139 (97.7%) |
| Seattle | | | | | | | | |
| Seattle | Vote 2000 | Not Registered in 1998 | | | Abstained in 1998 | | Voted in 1998 | |
| | | Not Registered for 2000 Primary | Abstained 2000 Primary | Voted 2000 Primary | Abstained 2000 Primary | Voted 2000 Primary | Abstained 2000 Primary | Voted 2000 Primary |
| Voted in 2000 General Election | NO | 33,990 (37.4%) | 25,779 (46.4%) | 1,088 (8.8%) | 283,860 (69.0%) | 3,284 (8.9%) | 35,609 (14.6%) | 8,401 (2.6%) |
| | YES | 56,976 (62.6%) | 29,791 (53.6%) | 11,311 (91.2%) | 127,666 (31.0%) | 33,764 (91.1%) | 208,877 (85.4%) | 317,901 (97.4%) |

TABLE 2: Forecasting Margin of Victory, Without Allocating Undecided RBS Respondents

| State-Office (total ballots) | Candidate | Party | Incumbent | Actual Vote ¹ | Registration Based Sampling | | | | | Random Digit Dialing | |
|------------------------------------|-----------------------|-------|-----------|--------------------------|-----------------------------|---|-----------------------------------|--------------------------------|--------------------------|----------------------|--|
| | | | | | Unweighted | Weighted to Original Strata Proportions | Weighted to All Registered Voters | Using "definitely vote" screen | Among Actual 2002 Voters | Unweighted | Using likely voter screen ³ |
| MD-Governor | Townsend | D | | 47.7% | 44.6% | 45.9% | 33.0% | 44.6% | 45.1% | 45.6% | 49.0% |
| | Ehrlich | R | | 51.6% | 48.0% | 45.3% | 42.6% | 48.4% | 47.6% | 49.4% | 49.4% |
| | <i>Margin</i> | | | <i>-3.9</i> | <i>-3.4</i> | <i>0.7</i> | <i>-9.6</i> | <i>-3.8</i> | <i>-2.5</i> | <i>-3.8</i> | <i>-0.4</i> |
| | <i>Forecast Error</i> | | | | <i>-0.5</i> | <i>-4.5</i> | <i>5.7</i> | <i>-0.1</i> | <i>-1.4</i> | <i>-0.1</i> | <i>-3.5</i> |
| PA-Governor | Rendell | D | | 53.4% | 48.6% | 50.1% | 55.3% | 49.8% | 49.8% | 51.4% | 54.0% |
| | Fisher | R | | 44.4% | 40.3% | 39.7% | 36.7% | 39.2% | 40.5% | 32.9% | 35.0% |
| | <i>Margin</i> | | | <i>9.0</i> | <i>8.3</i> | <i>10.3</i> | <i>18.6</i> | <i>10.6</i> | <i>9.3</i> | <i>18.5</i> | <i>19.0</i> |
| | <i>Forecast Error</i> | | | | <i>0.7</i> | <i>-1.3</i> | <i>-9.5</i> | <i>-1.6</i> | <i>-0.3</i> | <i>-9.4</i> | <i>-10.0</i> |
| NY-Governor ² | McCall | D | | 33.5% | 27.6% | 27.2% | 26.2% | 28.8% | 27.3% | 27.8% | 29.0% |
| | Pataki | R x | | 49.4% | 42.4% | 41.8% | 40.9% | 42.2% | 43.0% | 42.1% | 45.0% |
| | Golisano | I | | 14.3% | 18.5% | 20.0% | 21.7% | 18.8% | 17.9% | 17.0% | 14.0% |
| | <i>Margin</i> | | | <i>-15.9</i> | <i>-14.8</i> | <i>-14.6</i> | <i>-14.7</i> | <i>-13.4</i> | <i>-15.7</i> | <i>-14.3</i> | <i>-16.0</i> |
| | <i>Forecast Error</i> | | | | <i>-1.1</i> | <i>-1.3</i> | <i>-1.2</i> | <i>-2.5</i> | <i>-0.3</i> | <i>-1.6</i> | <i>0.1</i> |
| NY-Atty General ² | Spitzer | D x | | 66.4% | 64.2% | 65.1% | 67.2% | 65.7% | 65.2% | 62.8% | 68.0% |
| | Irrizarry | R | | 29.9% | 18.6% | 17.1% | 15.4% | 19.1% | 19.0% | 18.2% | 18.0% |
| | <i>Margin</i> | | | <i>36.5</i> | <i>45.6</i> | <i>48.0</i> | <i>51.7</i> | <i>46.6</i> | <i>46.3</i> | <i>44.6</i> | <i>50.0</i> |
| | <i>Forecast Error</i> | | | | <i>-9.0</i> | <i>-11.5</i> | <i>-15.2</i> | <i>-10.1</i> | <i>-9.7</i> | <i>-8.1</i> | <i>-13.5</i> |
| NY Comptroller ² | Hevesi | D | | 50.4% | 46.5% | 46.2% | 45.5% | 47.0% | 46.8% | 47.2% | 49.0% |
| | Faso | R | | 46.5% | 34.1% | 34.6% | 35.3% | 35.2% | 34.3% | 32.2% | 34.0% |
| | <i>Margin</i> | | | <i>3.9</i> | <i>12.4</i> | <i>11.7</i> | <i>10.1</i> | <i>11.7</i> | <i>12.5</i> | <i>15.0</i> | <i>15.0</i> |
| | <i>Forecast Error</i> | | | | <i>-8.5</i> | <i>-7.8</i> | <i>-6.2</i> | <i>-7.8</i> | <i>-8.6</i> | <i>-11.1</i> | <i>-11.1</i> |
| SD-Senate ⁴ | Johnson | D x | | 49.6% | 39.5% | 39.3% | 42.0% | 41.3% | 39.7% | 43.8% | 44.8% |
| | Thune | R | | 49.5% | 39.5% | 40.1% | 35.9% | 40.0% | 39.7% | 38.6% | 41.1% |
| | <i>Margin</i> | | | <i>0.2</i> | <i>0.0</i> | <i>-0.9</i> | <i>6.1</i> | <i>1.3</i> | <i>0.0</i> | <i>5.2</i> | <i>3.7</i> |
| | <i>Forecast Error</i> | | | | <i>0.2</i> | <i>1.0</i> | <i>-5.9</i> | <i>-1.1</i> | <i>0.2</i> | <i>-5.0</i> | <i>-3.5</i> |
| SD - Governor | Abbott | D | | 41.9% | 31.5% | 31.1% | 30.2% | 31.3% | 31.0% | 37.4% | 36.3% |
| | Rounds | R | | 56.8% | 49.3% | 50.4% | 44.2% | 51.1% | 50.1% | 44.3% | 48.4% |
| | <i>Margin</i> | | | <i>-14.8</i> | <i>-17.8</i> | <i>-19.4</i> | <i>-14.0</i> | <i>-19.8</i> | <i>-19.1</i> | <i>-6.9</i> | <i>-12.1</i> |
| | <i>Forecast Error</i> | | | | <i>3.0</i> | <i>4.5</i> | <i>-0.9</i> | <i>5.0</i> | <i>4.3</i> | <i>-7.9</i> | <i>-2.7</i> |
| SD - House | Herseith | D | | 45.6% | 38.6% | 38.5% | 36.7% | 40.0% | 39.4% | 43.4% | 43.1% |
| | Janklow | R | | 53.5% | 44.3% | 43.4% | 48.2% | 43.2% | 42.7% | 41.8% | 44.5% |
| | <i>Margin</i> | | | <i>-7.8</i> | <i>-5.7</i> | <i>-4.9</i> | <i>-11.5</i> | <i>-3.2</i> | <i>-3.3</i> | <i>1.6</i> | <i>-1.4</i> |
| | <i>Forecast Error</i> | | | | <i>-2.1</i> | <i>-2.9</i> | <i>3.6</i> | <i>-4.6</i> | <i>-4.5</i> | <i>-9.4</i> | <i>-6.4</i> |
| Average Absolute Error (All Races) | | | | | 3.1 | 4.4 | 6.0 | 4.1 | 3.6 | 6.6 | 6.4 |

Sample Sizes

| | Registered Only | | | | | | |
|----|-----------------|-----|-----|-----|-----|------|-----|
| MD | 838 | 838 | 838 | 657 | 590 | 960 | 725 |
| PA | 745 | 745 | 745 | 687 | 620 | 1214 | 636 |
| NY | 735 | 735 | 735 | 656 | 575 | 1018 | 624 |
| SD | 438 | 438 | 438 | 380 | 393 | 438 | 353 |

Notes:

- 1 - The Actual Vote reflects the percentage of the vote received by the major candidates. It may not sum to 100% because minor party candidates were excluded. The results were downloaded from the state elections office website in each state.
- 2 - Vote totals for candidates appearing under more than one party in New York have been combined. The party listed in the second column is the major party affiliation.
- 3 - This column reflects post-survey weighting of the data for vote likelihood and other factors as done by the Washington Post (MD), and the Quinnipiac Poll (PA & NY) for public release. An index of several questions about likelihood to vote was used in New York and Pennsylvania, while a single question about likelihood to vote ("How likely are you to vote...") was used in Maryland, and South Dakota.