

Ask the Expert: Polling the 2018 Elections

Scott Keeter*

Tags: election polling, telephone survey sampling methods

Survey Practice

Vol. 11, Issue 2, 2018

As the 2018 elections approach, polling finds itself in the spotlight once again and the glare is harsh. Facing the growing problems that confront all of survey research as well as public skepticism about polling that followed the 2016 presidential election, polling practitioners have examined their methods and many have made changes. For this edition of “Ask the Expert,” *Survey Practice* approached a wide range of public pollsters and sought their views about the environment for election polling this year and any changes they have made in how they conduct their polls.

The pollsters in this review all conduct public election polls for universities and major news organizations. Multiple pollsters using each type of methodology were contacted: online opt-in sampling, RDD phone, and RBS phone. Most of those contacted provided responses, but some declined or did not respond. Pollsters were asked to address the key recommendation from the [American Association for Public Opinion Research's \(AAPOR\) Ad Hoc Committee on 2016 Election Polling](#) that respondent education level should be included among their weighting variables, as well as issues related to sampling and identifying likely voters. They were also asked for comment on the state of the profession and the current environment for election polling. Responses have been lightly edited for clarity and style.

Of course, there are many other organizations conducting election polls than those represented here. Consequently, this should not be viewed as a scientific sample of current practice.

The pollsters interviewed for this article include the following:

Chris Anderson, co-founder and president of Anderson-Robbins Research, part of the bipartisan team (with Shaw & Company Research) that conducts polling for Fox News Channel

Rachel Bitecofer, Assistant Director of the Wason Center for Public Policy, Christopher Newport University

Scott Clement, Washington Post Polling Director

Charles Franklin, Director, Marquette University Law School Poll

Chris Jackson, Vice President at Ipsos

Ashley Koning, Director, Eagleton Center for Public Interest Polling, Rutgers University

Mileah Kromer, Director, Sarah T. Hughes Field Politics Center, Goucher College

Bill McInturff, partner and co-founder, Public Opinion Strategies and part of the bipartisan team (with Hart and Associates) that conducts the NBC News/Wall Street Journal Poll

Lee M. Miringoff, Director, Marist Institute for Public Opinion

Patrick Murray, Director Monmouth University Polling Institute

Doug Schwartz, Director, Quinnipiac University poll

Andrew E. Smith, Director, University of New Hampshire Survey Center

Harry Wilson, Director and **David Taylor**, Associate Director, Institute for Policy and Opinion Research, Roanoke College

FOR PHONE POLLS, WHO YA GONNA CALL: RDD OR RBS?

Most of the pollsters who responded are still using random digit dial (RDD) samples for national polling, but there is considerable diversity at the state level, with registered voter samples (RBS) in common use along with online opt-in methods. All of those using RDD stressed that they sample both landlines and cellphones and most have increased the share of cellphones they call. Doug Schwartz, director of the Quinnipiac University poll, gave a typical response:

... Cellphones have become more productive—in some cases just as productive, if not more, than landlines. With that in mind, I don't see the problem with simply continuing to raise the percentage of people that we reach on their cellphones as more and more people get rid of their landlines.

But others said that respondent cooperation on cellphones is declining. Charles Franklin of the Marquette University Law School poll wrote:

One new challenge for us is a decline in cell response rates over the past two years, seemingly related to a notable upturn in scam calls to cellphones in the state over that time. This, of course, drives up costs and drives down response rates.

One exception to the use of RDD at the national level is the NBC News-Wall Street Journal Poll (NBC/WSJ), conducted by a bipartisan team from Public Opinion Strategies (POS) and Hart Research Associates. Bill McInturff,

partner and co-founder of POS, speaking on behalf of his colleague Fred Yang at Hart and their teams said,

The most important change we made after 2016 was to switch to registered voter samples. We believed given what we were doing in our own firms that RBS offered some powerful advantages for the NBC/WSJ poll. Among those are that we now monitor the nonparticipants, looking for any evidence of nonresponse bias with respect to partisanship. The partisan scores in the sample allow us to estimate whether those who participate are different from those who don't. So far, we haven't seen evidence of this. But we know that it can happen. Our firm was polling in the Georgia 6th congressional district special election runoff in June 2017. Given the tremendous amount of money being spent in that race, voters were barraged with calls and other messages. We found that Republican voters, perhaps in response to the number of calls they were receiving, were less likely to respond to our poll, while Democratic voters seemed fired up and willing to talk. From this we concluded that the public polling might overstate the Democratic candidate's level of support, which turned out to be the case. So, the fact that we have not seen any partisan bias in response thus far in the NBC/WSJ poll doesn't mean it won't happen. We plan to continue to monitor this through the election.

At the state and congressional district level, RBS has been adopted by a growing number of pollsters. Its utility for congressional district polling is undisputed and the value of past voting history on the records is clear, even if how to use that information requires considerable judgment by pollsters.

Patrick Murray of Monmouth has experimented with RDD, RBS and hybrid samples and now thinks RBS is best, though not without its issues. One of the big advantages for him is what the information on the files makes possible in terms of analytics:

And because we have real data on the voters (including party registration, vote history, and even which precinct they live in) we can go back to the election results and see how well our model did. What we found is that there are slight differences in each of the 3 specials we polled. The Alabama Senate election had a clear surge of Democratic voters in Democratic areas, the Pennsylvania 18th House District looked closer to a typical midterm in terms of who turned up to vote and where, and the Ohio 12th House District was a hybrid of high turnout in Democratic areas against lower turnout in rural Trump areas which was offset by higher turnout in moderate GOP suburbs.

Murray does note one limitation of the RBS data, a problem also noted in a Pew Research Center (2018) [analysis](#) of voter files:

...[T]he one thing that some list providers are bad at is imputing education level. So, we tend to have to use self-reports and weight to an estimate drawn from Census data. But the vendors are good on some other measures (gender, race) based on the verification we do with comparing self-reports to the voter list values.

Other pollsters in this review who use RBS exclusively or nearly so include Harry Wilson and David Taylor at Roanoke College and Rachel Bitecofer at Christopher Newport University. Wilson wrote that "We moved from RDD to RBS in 2017, and I suspect we will continue with this type of sample." Bitecofer said, "We exclusively use voter files for our election polling..." and does not anticipate any change.

But not everyone is a fan of RBS. Doug Schwartz of Quinnipiac University put it very directly:

As for the way we sample and interview, I am a strong proponent of RDD sampling and using live interviewers, which are gold standard practices.... I think that the best response to these challenges [gauging who will vote] is to keep doing what we've been doing. Despite these challenges, I haven't seen any other methodology beat RDD in terms of accuracy, and that's what I value most as a pollster. I think this is especially true in years that the electorate may include a larger share of new voters or voters that don't typically vote, which 2018 may turn out to be. A great example of this is in the 2017 Virginia governor's race, where turnout did not look like the most recent governor's races. Our poll was the only one to call Northam's win exactly on the margin.

Lee Miringoff of Marist University wrote that they will continue to use RDD with dual frame samples and live interviewers. He argued that pollsters who used RBS in 2016 produced mixed results:

It revealed the limitations of RBS sampling particularly for state polling due to the inconsistency of state records, an inability to keep lists updated especially in states with more relaxed registration laws, and a literal and figurative disconnect when it comes to matching records with contact information.

Charles Franklin of the Marquette University Law School Poll also sees RDD as the best fit for polling in his state:

We are an RDD survey and don't plan to change for the

foreseeable future. We have increased the cell proportion each cycle and are now at 60% cell (vs 50% in 2016.) We do continue to consider alternative designs, including RBS, and for special samples of the Milwaukee area an ABS mail survey. We have also run comparisons with online samples from various sources but remain a bit concerned about coverage and size of available online panels in Wisconsin.

Franklin notes that Wisconsin is a high turnout state and that identifying likely voters has not been a major source of error in their polling. Consequently, he believes that some of the advantages of RBS in contributing data to likely voter models would be less consequential there.

In a similar vein, Andrew Smith of the University of New Hampshire noted that his state has characteristics that make it more amenable to RDD sampling than RBS:

In NH, we will still be using RDD (75% cell, 25% landline) because New Hampshire has significant migration as well as same day registration. In past elections, 10% to 15% of voters registered on Election Day. Consequently, RBS coverage is not very high in New Hampshire—about 65% last I checked.

However, his polling in a neighboring state, Rhode Island, is different:

We will be polling in Rhode Island this election and will be using an RBS sample. Rhode Island is a relatively non-mobile state that has had a very stable electorate, and there is no evidence of significant in- or out-migration in recent years. RBS coverage [of the likely electorate] is almost 80% because of the stability of the electorate, which gives us confidence that it is appropriate in Rhode Island, especially during what looks to be a low turnout midterm.

Mileah Kromer, of the Goucher College poll, said that the choice of RDD was "mission driven" by the endowment of the Sarah T. Hughes Field Politics Center, which she directs:

The Goucher poll is a policy poll for the purpose of representing the views of all Maryland residents and not just registered voters. By interviewing Maryland adults for the policy content and then screening for registered and likely voters for our election reports, we can both hold true to our mission and also do election polling.

A couple of pollsters had used hybrid RDD and RBS samples, but only one of those who participated in this discussion continues to do so. Scott Clement of the *Washington Post* wrote that their state-level surveys now include a mix of half RDD and half RBS interviews:

The *Post's* national polls continue to employ dual frame random digit dial samples with 65 percent of interviews completed on cellphones. For state-level surveys, we have adopted an approach where half of our sample comes via a registered voter list sample while the other half comes from random digit dialing; both sources include cellular and landline phones, and the samples are drawn such that a phone number can only be selected from one of the sources. The voter-list sample allows us to rely on validated turnout measures in past elections for a significant portion of our sample, while the RDD portion increases the coverage of our sample to include the portion of the voter-list sample that lacks a phone number. The voter-file portion of the sample also allows us to test the accuracy of our questions measuring likelihood to vote by looking at whether our respondents actually voted. We think this method is an effective tradeoff between the advantages of RDD and voter file samples and served us well in last year's Virginia gubernatorial election and the Alabama [U.S. Senate] special election.

By contrast, Patrick Murray of Monmouth University tried a hybrid sample approach but decided to return to RBS exclusively for his 2017 and 2018 efforts at the state and local level. The concrete nature of the RBS samples provides a more solid starting point for creating likely voter estimates:

I think those of us who do election polling really need to come to terms with the fact that it is impossible to conduct a true probability survey. We are trying to select a sample from an unknown population—that to me violates the premise of probability—how can each element have a known chance of being selected into the sample if we don't know which elements are actually in the population in question. (And quite frankly, I think a 5% response rate for a general pop RDD survey has a dubious claim to being a probability sample.) However, we do know what a voter pool (all RVs [registered voters], voters who have voted in the last two elections, and so on) should look like. And the voter lists provide reliable—if not always valid—data of what that population should look like. So, basically, we are going back to the old George Gallup sampling method (which quite frankly, I'm no longer convinced was a primary cause of the 1948 "failure")—and this is what the online pollsters are doing anyway. I just haven't moved to online because we lack the ability to do that using voter list data.

RBS sampling is at the heart of the most ambitious public polling [effort](#) of the cycle, which is being conducted by The New York Times Upshot and the Sienna College Research Institute. They are hoping to complete approximately

100 polls in competitive U.S. House races by Election Day. Nate Cohn, a domestic correspondent for The Upshot, described the choice of RBS and its rationale in an introductory [article](#) about their project:

In the abstract, there's a fine debate to be had about which approach [RDD or RBS] is superior. But we had no choice: It is hard or even impossible to conduct a poll of a congressional district with random digit dialing, since no lists exist of telephone numbers sorted by congressional district.

Calling voters from the voter file also has advantages, because it means we know a lot about our respondents before we've even spoken with them. In most states, we have their age, their gender, a reliable indicator of partisanship like party registration, and a record of when they voted (though not for whom). We know where they live, so we know whether their neighborhoods voted for Hillary Clinton or Donald J. Trump. We know their names, which along with their neighborhoods can give us a good idea of their race.

We can use this information in many ways—most important, to make sure our sample is representative. We know how many registered Democrats or Republicans live in the district, so we can make sure we have the right proportion of each in our poll. We can also use this data to model the likely electorate because past turnout is a pretty good predictor of future turnout.

APPROACHES TO SURVEY WEIGHTING

One of the principal findings of the AAPOR 2016 poll review was that some state-level surveys overrepresented college graduates and consequently overestimated support for Hillary Clinton.

Education was more strongly correlated with partisan vote choice in 2016 than in earlier years, which may explain why surveys that did not weight on education were nevertheless accurate in previous elections. While some pollsters contacted for this roundup have since changed their weighting to include education, others argued that this was unnecessary or inappropriate for their situation.

A clear example of a 2016 poll where weighting on education would have improved the estimates was in New Hampshire. Andrew Smith, director of the University of New Hampshire Survey Center wrote that

In our review of our 2016 polling, which significantly underestimated GOP support for Trump and other candidates, we found that we did not have enough people with "some college" education in our samples. When we included an education weight, our estimates matched the outcome of all

races. We have included an education weight in all of our polling since.

Ashley Koning, director of the Eagleton Center for Public Interest Polling at Rutgers University, also reported making this change:

We started regularly including education in our weighting in the summer of 2017, which was the first time since the election when we conducted our next big bout of public statewide polling in New Jersey. The move was of course in large part inspired by the election and the AAPOR task force report. Especially given New Jersey's incredibly diverse make-up as a state—almost like a microcosm of the entire country—we felt it was vital to start including education and we still do to this day. We now typically weight on gender, age, race, ethnicity, and education.

Another university pollster to make the change was Mileah Kromer of Goucher College:

I read the AAPOR report and discussed its findings with other academic pollsters and it was clear that this was an important change to make. Maryland is one of the best educated states in the country, but we certainly would not want to overstate education levels in our polls.

Lee Miringoff of Marist University was more skeptical about weighting by education:

I think the jury is still out. Although using an education weight improved polling estimates in some states, it did not in others. The recommendation to adjust to education for 2018 was based on analysis retrofitting pre-election survey data from 2016 to the final outcome. We will see in 2018 how that does or does not improve estimates.

This skepticism may be related to the fact that Marist already weights by income, which is correlated with education:

The Marist Poll has always weighted by income, a factor recently excluded by many other polling organizations because of increasing nonresponse. Instead of dropping income as a weight factor, we added follow-up questions to encourage a respondent to place themselves in a broader category which could still be used for weighting. Nonresponse is currently not any higher for income than for education in our surveys.

Miringoff noted that experiments they conducted with education in the weighting found no improvement over current practice in the quality of estimates.

Two pollsters active in Virginia reported that they did not plan to add education to their weighting protocol. Rachel Bitecofer of Christopher Newport University said:

I am not convinced that college educated voters are universally over weighted (and by extension non-college educated voters under weighted) in horserace election polling. I read the AAPOR report and analysis looking back at 2016 and applying the education weights suggested by the census data. I think it is quite possible that education weights, like most other weights, are variable among different populations. For instance, adopting the suggested weighting system for our final Virginia survey (which was accurate) makes the same survey inaccurate. I find a similar issue when I weight our final 2017 gubernatorial poll to the suggested education weights. The data goes from [Democratic gubernatorial candidate Ralph] Northam +6 to a dead heat and I should also point out that the one survey I am aware of that did use the education weights in Virginia erroneously predicted a very close margin. I will concede that Virginia may be an outlier. Certainly, after predicting two elections in Virginia correctly we are in an "ain't broke, don't fix it" approach. If something was to change in terms of our accuracy, I would certainly reassess our approach.

Similarly, Harry Wilson of Roanoke College wrote that

We have not included education among the variables we weight, and we would not have done that this year. Weighting for age, race, and gender has generally kept the education variable relatively stable in our election polling—in presidential election years, midterms, and in our gubernatorial election polling. I understand why some may choose to weight for education, but I think, again, we estimate that at our peril. Plus, our final 2016 poll had Clinton +7% and she won Virginia by 5.4%, so we felt pretty good about that! In fairness, we were not as close in the Virginia governor's race last year, but I think that was related more to turnout and party than it was to education. I can only wish that our polling was as good as hindsight!

Wilson's colleague David Taylor added:

We haven't found the need to weight for education in Virginia for state-wide elections, possibly due to the large effect of Northern Virginia and the Tidewater areas, which tend to be more educated and control more of the "vote" share in Virginia.

Scott Clement noted that education has been a part of the *Washington Post's* weighting protocol for many years but noted an additional change they made

in 2016:

...We refined these weights in the 2016 election by ensuring our samples matched not only the percentage of college graduates in the population but also the share with post-graduates/bachelor's degrees, as these groups have proven to be politically distinct. Education is clearly one of the most important factors in views toward President Trump, so we will continue to ensure our samples match the population, in addition to the other demographics we weight toward.

One place where education weighting did not make a difference in 2016 was Wisconsin. Charles Franklin of Marquette University observed that

We have always weighted to education, which did not keep me from missing the 2016 race in Wisconsin which we had +6 for Clinton vs the +1 for Trump outcome. We continue to use education in our weights, along with sex, marital status, age and region of the state. We don't weight to race because there is a modest non-white or Hispanic population and our weighted data come quite close to those population values without explicit weighting.

That raises the question of what did cause the Marquette poll to miss the mark. Franklin noted that the underestimation of Trump was concentrated in the Milwaukee suburbs and north toward the Green Bay area.

These well-educated, professional, and very Republican areas were the problematic area for us, rather than the "white working class" or rural areas of the state where Trump did well but our poll picked that up. We are doing a "leaned vote" now [asking those who declined to state a candidate preference which way they were leaning], in response to that error. We did not have a large undecided in 2016 but perhaps a stronger "leaned vote" might have picked up some of the apparent ambivalence about Trump in GOP strongholds that was then resolved into votes for him.

Nate Cohn of The Upshot [notes](#) that while weighting by education is highly desirable, it is not so easy to implement with RBS because of the absence of a reliable measure on the voter file. Instead of using the education measure on the file, they create targets using an amalgam of Current Population (CPS) and American Community Survey (ACS) estimates, adjusted for what Cohn calls

...the consistent evidence of a turnout surge among well-educated voters. We use the validated turnout in Upshot/Siena poll data from Virginia and Ohio 12 [House District special election] to try to tease out how much the turnout among

well-educated voters has increased since 2014. The effect of this adjustment is to increase the college-educated share of the electorate by two to three percentage points over a 2014-based model.

Chris Anderson, a Democratic pollster who is part of the bipartisan team that conducts polls for the Fox News Channel, employs a similar approach for similar reasons:

We identified the importance of utilizing a 4-point education weight in early 2016 which, we believe, was a key reason our polls were accurate. However, estimating the appropriate education profile to weight to continues to be a challenge. Estimates of education level at the individual voter level tend to include a lot of missing data, with most voter files missing education estimates for a third of voters or more. Exit poll estimates are suboptimal, having been documented to skew more educated. As a result, we are generally weighting to ACS estimates for education, while taking into account differential registration and turnout among voters of different levels of education.

CHALLENGES IN 2018

Asked about the challenges they are facing in 2018, the pollsters responded with answers about the growing difficulties in contacting respondents and gaining cooperation as well as broader issues related to the political climate.

Several noted that with the political environment favoring Democrats this year, gauging turnout may be a special challenge. Past midterm turnout may not be as predictive as it was in 2010 or 2014. Accordingly, likely voter modeling is a special focus, though some pollsters indicated that they did not like to share details about this aspect of their work.

Chris Anderson put the challenge this way:

...There seem to be more unknowns and questions about who will and won't turn out than in previous cycles. However, we view this as more of a matter of degree of uncertainty as opposed to a new dynamic. Our approach has always been to build as few assumptions as possible into our likely voter modeling. Specifically, for state polls when we use voter file samples, we use a broad vote history [to] select and then a relatively light likely voter screen. This isn't a new practice for us; however, it's one that seems more important than ever.

Lee Miringoff is critical of the trend toward the imposition of assumptions in likely voter modeling but noted that the trend is understandable:

What had made public polling unique was its goal of

interviewing everyone and identifying reliable and valid measures to understand who would vote and why. It didn't start with past assumptions. Today, in an effort to shore up limitations of nonresponse, high costs, and improve speed, there is an emphasis on identifying the "right model" of the electorate.

Charles Franklin and Mileah Kromer are also in the simpler-is-better camp: Franklin wrote:

We have always used a simple likely voter screen based on self-reported certainty of voting (in a sample of self-reported registered voters only). This has worked well for us in estimating turnout within region of the state and overall. Wisconsin is a high turnout state, so a very large proportion of registered voters actually turn out in presidential years, though obviously less so in midterms. In comparing our self-report filter with model-based turnout estimates, I've found very modest differences in classification with no consistent advantage to either method.

And Kromer responded:

Our likely voter model doesn't use past vote. We rely on respondent answers about interest in the election and intention to vote—the top two points on the scale ("certain" or "very likely").

Harry Wilson described the hazards of imposing one's assumptions about changing conditions on the models. In 2012, he did not believe that Obama would be able to fully recreate his 2008 coalition in Virginia. Accordingly:

I modified our weighting standards ever so slightly from the 2008 exit poll to reflect a slightly more Republican, older, and whiter electorate. Vote counts proved me wrong. Our last poll would have been "off" anyway, but I increased that margin by 1% or 2% by tinkering with the weights. I was later quoted (accurately) by a reporter who inquired about why several polls, including ours, were incorrect, and I said that "I drank the Republican Kool-Aid." I learned from that experience, and I try not to substitute my judgement for the last exit poll. One is concrete, and the other is not. I also do not sip Kool-Aid of any sort!

Andrew Smith offered a similar perspective:

... I've learned over the years that my estimations of what the electorate will look like are not reliable, nor do I think anyone else's are either. My position is that we have to rely on the respondent to tell us if they will vote, not arbitrarily decide if they will by using a scale or cutoff system. What we do to try to get

around the social desirability bias is to make it as easy as possible for a potential respondent to say that they will not vote. We ask registration status, interest in the election, how much they talk about the election, and then ask how likely they are to vote. If they say they will definitely vote or vote unless some emergency comes up, we include them.

COMMUNICATING RESULTS TO THE PUBLIC

Pollsters commented on problems related to the communication of poll results—and especially uncertainty—to the public. Scott Clement noted that even the most familiar indicator of uncertainty, the margin of sampling error, is not well understood by audiences:

Polls are often assumed to be more precise than they are, in part because the margin of sampling error is not taken into consideration as much as it should. This can be a challenge in pre-election polls, because the margin of sampling error applies to each candidate's support, and a candidate's lead usually needs to be 1.5 to 2 times the size of the error margin to be considered statistically significant.

Lee Miringoff suggested that a major problem in 2016 was not the polling but the forecast models that communicated too much certainty about the outcome:

As a member of the AAPOR committee to evaluate the 2016 election, I was struck more by the misdirection of the election narrative powered by the forecasting models than by mishaps in methods of "gold standard" polls. Scientific national and state polls conducted close to the election were mostly accurate. The "shock" of 2016 was more a result of the forecasting models, a cottage industry for 2016, which all pointed with a high degree of certainty that Clinton would be the likely winner.

Patrick Murray of Monmouth University is one of a few pollsters attempting to communicate uncertainty by showing multiple likely voter models. He decided to do this on the basis of observing that pollsters' judgments about turnout tend to be flawed. Lee Miringoff expressed a similar view, "Pollsters looking forward to the election contests of 2018 must not be so eager to adjust to what in retrospect would have 'fixed' the problems of 2016. The industry should avoid fighting the last war."

In this vein, Patrick Murray described several adjustments he is making in both sampling and weighting but concludes:

While I am confident in this approach we still have to admit that there is a lot of room for uncertainty. Which is why I am committed to showing multiple likely voter models in our

polls—first to underscore the uncertainty and second to demonstrate that pollsters can make different, but defensible, decisions on what constitutes a likely voter that can lead to very different estimates of the electorate. And we, especially the media, need to do a much better job coming to grips with the nature of the beast. We are not in the fortune-telling business.

Chris Jackson of Ipsos also noted that pollsters can do better at communicating uncertainty. He said that the Ipsos-Reuters poll would show multiple alternative likely voter models in their releases. In their project with the University of Virginia's Center for Politics, they plan to broaden the range of inputs beyond polling:

... Best practice in forecasting understands that no single approach leads to knowing the future. The most robust predictions use multiple, independent indicators and looks for common conclusions or divergent directions. A key failing in 2016 was the overreliance of polling as the primary or sole input into forecast models. Our partnership with the University of Virginia looks to broaden the information used as can be seen on our Political Atlas. Here we are displaying a multilevel regression with poststratification model of our poll, with expert opinions from the Crystal Ball [a nonpartisan political analysis and handicapping newsletter at the University of Virginia], along with sentiment tracking on social media. As we close to the election, we will also layer in structural models. By showing all these different, and potentially differing, types of forecasts, we are more honest with readers about the certainty we have in a particular election outcome.

A similarly broad scope is being pursued in the real-time midterm U.S. House election polling being conducted by the Upshot/Sienna. In addition to the novel feature of showing the results updating as calls are completed, the findings are shown for several different turnout models and four different weighting schemes. For each race being polled, the research team flags what they believe to be the best estimate but acknowledge that different assumptions about turnout or different approaches to weighting could produce different results.

CHALLENGES FACING POLLING IN THE AGE OF PRESIDENT TRUMP

After a period of relative stability in response rates, public opinion polling once again began to experience greater noncontact and lower cooperation rates after 2016. The cause of these changes is some mix of the secular trends that have been hammering the industry for the past few decades and newer issues that stem from the highly charged political environment of today.

Asked about the specific challenges facing polling in the era of President

Trump, Ashley Koning of Rutgers wrote that maintaining the public's trust in the process and the science behind it is more difficult now:

The number of times we've had to reassure a client that the process still works and the number of skeptical questions I get when giving talks has certainly increased since 2016. I think this new era of politics has brought about a significant amount of distrust, and on top of known challenges (like response rates) that we have already been dealing with in recent years, this distrust further hampers our job to provide an accurate depiction of public sentiment.

But to Koning, trust is not the only issue:

Another challenge is fear—the fear to answer certain questions, the fear of repercussions, the fear of any loss of confidentiality or anonymity. How do we represent the public when the public doesn't want to answer us ... when the public is too *afraid* to answer us? While this debate is at the forefront with the Census' citizenship question, we faced these challenges in the summer of 2016 in New Brunswick, New Jersey. With the 2016 presidential campaign at full speed coinciding with local ICE (Immigration and Customs Enforcement) raids in the city during our fielding period, respondents showed a much greater reluctance than in the past to respond to our community survey, were less likely to identify as Hispanic (when all indicators said otherwise), and less likely to take the survey in Spanish. It was stark evidence of what the new political climate could do to the voice of populations that were already underserved. As a strong believer that polls are a great equalizer and one of the most representative participatory acts, the impact that fear can have on the polling process is one of the most daunting issues of all in this era.

To Charles Franklin of Marquette University, heightened political polarization has also had an impact on what pollsters are trying to measure:

I think the extreme polarization of politics, nationally and specifically in Wisconsin, poses some challenges. For example, I now generally leave out mention of the president or governor in asking about issues because doing so seems to "snap" issue positions into partisan camps. Likewise, the dynamics of votes or issues is pretty constrained when 90% or more of partisans are reliably with their group, leaving it to independents to provide almost all movement. That is the world as it is, so, of course, we want to capture it, but it constrains results quite a bit.

Another consequence of polarization has been heightened concerns among potential respondents about the possibility of sponsor bias. Bill McInturff

noted that:

Historically, we have introduced the poll as being conducted on behalf of NBC News and the *Wall Street Journal*. After 2016, we stopped doing that so as to minimize the chance that we'd lose respondents who might not trust mainstream news organizations.

McInturff also noted that, in a nod to growing resistance from respondents, they have reduced the average length of their poll from 20 minutes to less than 15 minutes.

The current era is bringing challenges to interviewers as well. Harry Wilson of Roanoke College observed:

I think our biggest challenges in polling in the Trump era have been keeping up the morale of our interviewers and keeping interviews short. People/respondents are angry, and unfortunately our interviewers can get the brunt of that...from both sides. That also tends to make the interviews run longer, and we must work harder to get interviewers to steer folks back to the questions, especially if they get off course several times.

For Lee Miringoff, the rising costs and declining cooperation have brought about considerable methodological experimentation, but this has also brought about reputational risks for the entire profession:

Experimentation is invaluable and required if the survey profession is to survive. Unfortunately, instead of occurring in the "lab," new methods are broadcast real-time with little vetting except for wins and losses in the post-mortems. Even for the most commendable of efforts, real-time testing feeds the narrative of "fake news" and "fake polls" when survey results from these "experiments" miss the mark.

Andrew Smith of the University of New Hampshire believes that the bigger issue faced by pollsters is that there may not be much there to measure: that most people do not know very much and are not particularly interested in politics:

We spend our time focusing on methodological minutiae and ignore this much bigger problem. This can be addressed if we presume don't know and don't care is a valid expression of public opinion, and not try to force people into fitting into our categories. We also should encourage people to say don't know/don't care as this, in my mind, is closer to the truth. Most people care about their day to day lives, not political and policy issues. And the growth of the Internet, TV, Netflix, etc. makes it

even easier for people to ignore politics/policy than they did 30 years ago when most people were still pretty much stuck seeing the news or getting a newspaper. We should review Anthony Downs and take seriously his observation that voting is a non-rational economic activity—our individual vote really doesn't matter! It seems that much of America has figured this out!

It is clear from this review of practitioners that many pollsters took heed of the implications of key findings in the AAPOR 2016 polling review, most notably the importance of including education among the set of weighting parameters. But some said that weighting by education was not appropriate for their organization. Another finding in the AAPOR report is that Trump supporters were more numerous than Clinton supporters among voters who revealed their preference after the election. At least one pollster in this review is adding leaned vote to his approach but none volunteered that they planned to poll closer to Election Day in order to capture more late decisions. In fact, Charles Franklin noted that they had a policy not to release a poll later than Wednesday of the week before Election Day so as not to dominate the news too close the election.

The AAPOR review said that errors in identifying likely voters was a possible source of error in 2016 but that the data available at the time of the review was inconclusive. As this roundup has shown, pollsters continue to have a wide variety of opinions on how best to gauge who will vote. But considerations on this issue have driven the most striking methodological trend among public election pollsters in recent years, which is the growing adoption of RBS sampling. Still, many pollsters remain fans of RDD and offer strong reasons for their decision not to change.