

# **A PRIMER ON PRE-ELECTION POLLS:**

Or

*Why Different Election Polls Sometimes Have Different Results*

**Cliff Zukin,**

Professor of Public Policy and Political Science

Bloustein School of Planning and Policy & Eagleton Institute of Politics,  
Rutgers University

The views expressed here are solely those of the author

Primary Election Version

December, 2015

*Every major campaign cycle brings with it a new wave of pre-election polls and a new round of concerns and criticisms about the outcomes of those polls. And as the number and types of polls increase, so does confusion about how to evaluate different poll results. It is here that AAPOR can lend a hand to journalists, researchers, policymakers, and the general public by explaining election polling and provide the tools needed to make judgments about the plethora of polls out there. To support this knowledge sharing, we have asked former AAPOR President Cliff Zukin to update his primer on election polling. This primer is not meant to represent AAPOR's formal position, and some may disagree with some elements here, but it provides a great background for those who are interested in understanding the complexities of this ever-evolving field. We hope that you find answers to your questions in this primer and we sincerely thank Dr. Zukin for this excellent contribution. For more AAPOR resources on election polling, please see <https://www.aapor.org/AAPORKentico/Education-Resources/Election-Polling-Resources.aspx>*

*Mollyann Brodie, 2015-2016 President, AAPOR*

## **INTRODUCTION**

Election polls are a special breed among public opinion surveys. They call for more judgments—the art rather than the science of the craft—on the part of the pollster than other types of polls. And this brings into play a host of other reasons why the estimates of well-established and well done pre-election polls may differ from one another, even when these polls are conducted at a similar point in time. Also, the polling landscape has changed dramatically, even as recently as two presidential elections ago in 2008. Because of the growth of cell phones, declining response rates and increasing difficulty

in determining “likely voters,” I have argued [here](#) that election polling has gotten harder to do well, and that 2016 will present election polls with a very challenging environment.<sup>i</sup> This primer is meant to be a guide for journalists, academics, and anyone closely following polls in the 2016 election.

## **SAMPLING AND REAL SAMPLING ERROR**

*Sampling is the foundation of scientific survey research* and is based on the branch of mathematics having to do with probability theory. Statisticians make a primary distinction between two types of samples. Probability samples are based on everyone having a known chance of being included in the sample. This is what allows us to use mathematical properties such as the Central Limits Theorem and Law of Large Numbers to be able to generalize from our *sample* back to the larger *population* from which it was drawn with knowable boundaries of how accurate the sample estimate is likely to be.

It is only appropriate to attach a margin of *sampling error* figure to probability samples. This figure, familiar to most journalists and polling observers, is expressed as “plus or minus” a number of percentage points<sup>ii</sup>, is the most commonly known source of variation for why polls may differ. In this election year, we often hear statements, such as prospective Democratic candidate leads the prospective Republican by three points, 47% to 44%, with a sampling error of plus or minus three percentage points.

What is less commonly known is that the margin of sampling error does not apply to the spread between the two candidates, but to the percentage point estimates themselves. If applied to the three point spread, the three-point margin of error would seem to say that the Democrat’s lead might be as large as six ( $3 + 3$ ), or as little as zero ( $3 - 3$ ). But when correctly applied to the percentage point estimates for the candidates, the Democrat’s support could be between 50% and 44 % ( $47 \pm 3$ ), and the Republican’s between 41% and 47% ( $44 \pm 3$ ). Thus, the range between the candidates could be from Democrat having a 9 point lead ( $50\% - 41\%$ ) to the Republican having a 3 point advantage ( $44\% - 47\%$ ). So, the sampling error is generally much larger than it may seem and is one of the major reasons why polls may differ, even when conducted around the same time.

In the primary election period, the sampling error is even more of a consideration, since sampling error is strongly related to the size of the sample. In a sample of 1,000 correctly selected respondents sampling error is about  $\pm 3$  percentage points, but a sample of that size would likely contain about 350 Republicans. On this respondent base of only Republican voters, the sampling error would be about 5 percentage points, so if candidate A had 27% he might actually have between 22% and 32% while candidate B at 19% might be between 14% and 24%.<sup>iii</sup> For candidates bunched back later in the pack, receiving less than 5%, polls cannot reliably be used to estimate their relative positions.

**MODES OF SURVEYS:** *Surveys may be conducted by telephone, Internet, mail or in-person. For all intents and purposes, election surveys have largely been conducted by telephone in the past, although more on-line surveys were used in 2014, a trend that is likely to continue in 2016.*

#### *Telephone Surveys*

Most pre-election surveys are conducted by telephone, using one of two types of sampling frames, or definitions of who is eligible to participate. The most common approach in the U.S. is what is called an **RDD** sample, short for random digit dialing. In this case samples of both landline and cell phone telephone area codes and exchanges are taken, and then random digits added to the end to create 10 digit phone numbers. The first step ensures proper distribution of phone numbers by geography; the final step, adding the random numbers, makes sure that even unlisted numbers are included. In the case of the landline RDD sample, a randomly designated respondent is then asked to participate in the survey.<sup>iv</sup> This is the standard practiced by most public pollsters.<sup>v</sup>

An alternative is called registration based sampling, or **RBS**, and this is the method used by most commercial pollsters working for political candidates and parties. This begins with a sample of individuals drawn from publicly accessible lists of registered voters, to which phone numbers are then matched. This is less costly and more efficient, as almost all calls result in reaching a working phone number and a registered voter, which is not true of an RDD sample. The primary disadvantages of RBS surveys are that they miss people who have recently moved or are *newly registered to vote*, which may be a non-trivial portion of the electorate in some states or the country as a whole. In fact, a great danger here is that these samples miss voters mobilized by specific campaigns, who may be more likely to vote in primary and caucus states.

Moreover, not all sampled through voter lists may be findable by telephone. Certainly there will be a bias towards those with **listed** telephone numbers,<sup>vi</sup> meaning a substantial portion of the electorate could be missed. Also, the purging and updating of voter registration lists historically have varied from state to state, so the accuracy of RBS sampling will vary, although official state lists have become more consistent in recent years. Because of the difficulty in finding likely voters in low turnout primary contests, discussed later, a number of polling organizations will be employing RBS samples in the early nominating contests.

#### *Cell Phones*

As recently as 2008 the exclusion of cell phones from telephone polls was common; just 10 years ago only about 6% of the country was “cell phone only.” Together, those using a cell phone exclusively or “mostly” now comprise 60% of the adult public.<sup>vii</sup> Today it is *impossible to do a credible telephone survey without including cell phones*. At the end of 2014, 44% of the adult population was cell phone only, while just 8% were landline only.

Federal law prohibits using automatic dialers to contact cell phones, so cell phone surveys cost much more than landline interviews to conduct. Some survey organizations have a tendency to skimp on the number of cell phone interviews. The very best survey organizations commonly include a mix of 50%/50% landline/cell interviews, and some are moving to 60% cell phones as usage continues to grow. Given the impact on coverage, the proportion of cell phones in the sample is one indicator of survey quality.

### *IVR Surveys—Robo Calling*

Some pre-election telephone polls use no live interviewers at all and rely exclusively on recorded voices. These are called **IVR polls**, for *interactive voice response*, or sometimes robopolls in the vernacular. Here an automatic dialer calls households and a pre-recorded voice asks questions and asks respondents to enter the number that corresponds with the survey response options given them. Because IVR polls are unable to dial cell phone numbers unless they are hand-dialed and added to their landline samples, they are unable to accurately represent the voting public. Some do add cell phones dialed separately from those dialed automatically, but some add only a trivial amount, largely for cosmetic purposes than for truly increased coverage. Again, the number of cell phones added to landline interviews is a fair an indicator of quality; if cell phones do not make up at least one-third of a telephone sample, it may not be worth reporting.

IVR polls can be problematic even when cell surveys are added to the mix, given that there is no respondent selection procedure within the contacted landline household.<sup>viii</sup> Given who tends to answer landline telephones, IVR polls tend to be disproportionately composed older women, who are likely to have a different response pattern and may have a particularized response to the two polled candidates than a cross-section of the public as a whole. Such polls use no selection technique when contacting a household, but instead try to compensate for the lack of within-household selection by weighting the resulting data to demographic targets. Not having a random selection within the household compromises a fundamental tenet of probability sampling and should require an accounting and justification before being reported.<sup>ix</sup>

### *Internet/On-line Polls*

The cost has driven many election pollsters to the Internet, where costs are a fraction of those required to do a good telephone survey. There are two major problems with this. First, is what pollsters call “coverage error”—not everybody has Internet capability. In 2014, about 85% of the public has Internet availability with about 70% having some sort of broadband coverage. This turns out to be a more severe bias in election polling. For example, while almost everyone (97%) under 30 uses the Internet, they made up only 13% of the electorate in 2014. In contrast, 40% of those 65 and older do not use the Internet, and they comprised 22% of voters in 2014.<sup>x</sup>

Second, pollsters have not yet figured out how to draw a representative sample of Internet users in the U.S., so most Internet surveys are based on non-probability samples. As noted, a probability sample is necessary for statistically valid inference to be made from the sample back to the population from which it was drawn, and a margin of sampling error computed.<sup>xi</sup> This assumption is not warranted in the case of non-probability samples. Although the large majority of Internet samples are non-probability samples, there are a few probability-based Internet samples, including some that address the coverage issue by providing Internet access to participants.<sup>xii</sup> *It is important for journalists and others evaluating polls to understand that any poll comprised of self-selected respondents, including call-in polls and Internet or Web surveys where people volunteer to participate in response to an open invitation are **non-probability** samples, the results of which should be interpreted with caution.* The vast majority of reputable public opinion polls continue to rely on probability sampling.

### **Timing and Field Procedures**

The timing, of course, refers to when a poll is conducted. And as pollsters are fond of saying, even the best pre-election poll is a no more than a snapshot in time. Polls don't predict; they describe the situation of the moment. Obviously, pre-election polls with different field dates may yield different results as voter preferences may change with time. However, a largely invisible reason for differences is that polling organizations have different field procedures. Field procedures refer to the ground rules under which the interviewing is done. And, there are some tradeoffs that are made by polling organizations.

For example, a field period of seven days would allow for a number of callback attempts to reach the selected respondent before allowing substitution with a new household and respondent. But campaign events may happen in those seven days, making that poll harder to interpret. A three-day poll may focus more narrowly on a particular point in time, but likely at the sacrifice of callback attempts. Callbacks matter because one respondent may not be the same as the next; extra field time may be necessary to reach younger voters, for example, who may be more Democratic in orientation than others. So factors like the number of callbacks, days of interviewing, and response rates may also be reasons why polls purporting to measure the same thing give different results.

*Tracking polls*, or polls where interviews are conducted every day either released on their own or aggregated with others to some consistent base, such as "the last 3 days," are a special case. Given that their callback procedure to reach their primary respondents is compromised by the length of time in the field, they may be more useful for spotting trends of voters moving up or down (reliability), and somewhat weaker at estimating vote choice (validity).

## Question Ordering and Wording

It has long been known that the ordering and wording of questions in a survey can affect the results. In ordering, responses to questions asked early in the interview may affect later ones, as frames of reference are set, respondents are conditioned to think about some topics more than others, and respondents attempt to be consistent in their responses to interviewers. For example, a survey that asked respondents a set of questions on the dysfunction of government before asking for whom they planned to vote could lead to a bias in favor of Trump or Carson, outsider candidates; a line of questioning on the willingness to vote for a woman as president could lead to an overstatement of intentions to vote for Clinton or Fiorina in subsequent questions.

In order to minimize this problem, most researchers will ask the horserace question (*If the election were held today, for whom would you vote...*) before any other substantive election question on the survey. (This does not include neutral questions about whether people are registered, or how interested they are.) After all, when people go into the voting booth they will have had no warm up questioning on issues or candidate qualities. However, perhaps in hopes of simulating the campaign, some polling organizations begin their surveys with substantive policy or election-related questions before asking about vote intentions. When interpreting poll results it is always useful to know the context in which a question was asked. While two polls may have asked the horserace question in the same form, one may have done so after unconsciously pushing some respondents in one direction or the other by earlier questioning. So question ordering also becomes a source of possible variation in the results among published polls. Best practice is to ask questions that might influence vote choice **after** the horserace question.

The wording of questions—even the horserace question—may also vary from one poll to the next. In the general election, some polls will ask a two-way vote intention question, naming only the major party candidates but recording all answers, while others will explicitly add a third party candidate's name. While most polling organizations asking about the candidates add their party labels as a cue, some may just name the candidates. And later in the season, trial heat questions that also name the vice presidential candidates may produce somewhat different results than when only the presidential candidates are mentioned. There is even some evidence that there is a slight bias in favor of whichever candidate is named first in survey responses, so some organizations **rotate** the names of the two candidates while other do not. So differences in question wording may also be a reason why polls have differences in their reported findings. In the early Republican polls, some polling organizations asked the question by naming all the (14-17) candidates; others asked it as an open-ended question, advantaging the better-known candidates.

## POST-SURVEY ADJUSTMENTS

*After the data are collected, researchers are called upon to make two adjustments to correct for any shortcomings, or biases, in who was interviewed, and to make a model about who is likely to actually vote in an election. These adjustments—weighting, and determining “likely voters”—constitute more of the art than the science of poll-taking, and account for a great deal of variations between the estimates of one survey and the next.*

### *Weighting*

Weighting is an important and common practice in survey research. Even the best polls cannot interview a perfect sample, due to non-response and non-coverage, among a variety of reasons. (Non-response occurs when people who are sampled refuse to take part in the survey or are never contacted during the field period; non-coverage occurs when not all people who will be voting are included in the sampling frame—an Internet survey would miss voters who do not have access to or use the Internet; a telephone survey would miss those without any phone service, for example.)

Weighting is a fairly straightforward task in surveys of the general public. Thanks to the U.S. Census we know how many people in the entire U.S. have a few fixed characteristics, such as age, education and race/ethnicity. When we look at who we actually interview in our samplings, we can adjust—or weight—for these characteristics to make sure they are correctly represented. For example, if we knew that **30%** of the adult **population** had graduated from a four-year college, and 45% in our **sample** of 1,000 graduated from a four-year college, we would need to weight these respondents by a factor of .67 to make sure the data reflect the correct proportion of this group. [The math here is straightforward: Take the 45% of college graduates, count each as .67 of a person, and they will contribute to the pool of all answers as if they were 30% of the total. (.45 x .67 = .30).]

As necessary, ubiquitous and powerful as weighting is in survey research, it is important to note important limitations of weighting as applied to election polling--the only population parameters researchers can have confidence weighing to must be (1) known and (2) stable. Neither of these is true in election polling, so there is much more guesswork involved. If there isn't a fixed known parameter in the *electorate* (as opposed to the population as a whole), such as the education level of voters, it can't be weighted to. And obviously, *one cannot weight to an event that has not yet happened*, such as what turnout will be among African American voters or conservatives or those over 50 years of age, for example.

Most pollsters of published surveys first ask a sample of the general population about their race, gender, age, etc., and then weight their data to what a random sample of the population should look like, and then go on to pull likely and non-likely voters out of that big (already weighted) general population sample. Some, however, weight to a

picture of what they *believe* turnout will be, based on past experience and elections -- and not everyone doing so is painting the same portrait (this practice is more common among campaign-sponsored polls than published polls).

This later strategy involves a fair amount of guesswork. In the 2012 election, the national exit poll estimated that African Americans made up about 13% of the electorate, and 93% of these votes went to Barack Obama. What if 9% of a polling organization's sample is made up of African Americans in 2016? What if it is 16%? It will obviously make a difference in the horserace estimate, but we won't know which is correct until Election Day.<sup>xiii</sup> And, of course, the past is no guarantee of the future. So, the pollster's dilemma is "What do we weight to?"

A second issue concerns whether polls should be weighted to reflect an assumed distribution of the electorate by the political party of respondents. A party identification question generally placed near the end of the survey, asks people to state whether they consider themselves to be a Democrat, Republican, independent or something else. The vast majority of pollsters do **not** feel it is appropriate to weight by party. The scholarly literature comes down firmly on the side that party is not a fixed attribute, like race or gender or age. It is an attitude, and peoples' responses to this question change based on circumstances and events. And indeed, the American public does show fluctuation in partisanship over time, as well as individual changes. A small number of pollsters do weight by party, but that is tantamount to guessing what the electorate will look like on Election Day, which of course is unknown.

Party ID is the most critical variable predicting the vote. In November of 2012, the national exit poll estimated that 92% of Democrats voted for Obama and 93% of Republicans voted for Romney. A two or three point difference in the estimation of the partisan makeup of the electorate will easily lead to polls that differ by two or three points, all other things being equal.

### *Likely (Probable) Voters*

Another, fundamental problem all election pollsters face is that there is an *over-reporting of the intention to vote*. When respondents' self-report of intentions in pre-election polls have been compared to actual turnout (again, known only after the election) we have historically found a large over-report of voting intentions. So the pollsters' dilemma here is to separate the wheat from the chaff: Of all those saying they will vote on Election Day, which ones will really do it, and which ones will stay home? And, of course, people change in their commitment to voting as the campaign unfolds. Respondents are probably better able to tell if they really are going to vote as it gets closer to Election Day. This means that the definition of a likely voter is somewhat of a moving target, compared to the definition of registered voters, for example.

Research finds no magic bullet question or set of questions that can reliably determine likely voters with 100% accuracy. Thus, different organizations have different ways of estimating who are probable voters. Most polls ask a combination of questions that cover three areas that are highly correlated with voting: a) self-reported vote intention; b) measures of engagement (following the election closely, interest, care who wins); and c) past voting behavior (voted in prior elections). They then combine responses to create an index that gives each respondent a total score. Most then use a cutoff point so that only the candidate preferences of the “most” likely voters are used, and the choices of others are discarded. But even while most use such a scale, the component questions that go into the scale differ, and so this too is a source of variation among polling organizations’ results.

There are other approaches as well. A more efficient but perhaps more risky strategy uses a single question or two of reported intention and does not complete the interview with those not passing the screen. For example, a poll might ask about someone’s chances of voting on a scale of 0 to 10, and only continue interviews with those who gave themselves a 10. This could result in only 40% passing this screen. If real turnout on Election Day was 56%, the underestimation might be biased if, say, Democrats were less likely to be among the initial 40% but overrepresented in the next 16% slice of the electorate. While most polling organizations use a cutoff point for likely voters (take all of those in the top 56% and none of those in the bottom 44%), others may give voters weights based on the probability of voting to everyone in the sample rather than using a cutoff. And still others may simply use a fixed set of screening questions that have worked well for them in the past, leaving a lot of room for variation in the vote choice estimates produced by different polls.

A second issue in determining likely voters is estimating *how many* there will be, which may affect the division of the vote. In 2012, the turnout of eligible voters was 58.6%; in 2008 it was 62.2%. What will it be in 2016? Should pollsters predict it will be stable, bump up to 2008 levels, or continue to decline? Suppose a choice of a cutoff point of 62% gives an estimate that the Democrat leads the Republican by three percentage points. But when pruning the expected electorate to 56%, it may be that the Republican leads the Democrat by 2 points. So, another source of possible differences is what percentage of voters is let in during the likely voters scoring process. Thus, while all polling organizations will release figures for who they believe are *likely* voters, no two organization will define them in exactly the same way. It is worth noting that estimates of likely voters generally come into congruence as the election gets closer.

Indeed, an analysis presented by the Gallup Organization at the 2015 conference of the American Association of Public Opinion researchers suggest that different estimates of turnout were a primary reason some polling organizations did not perform well in signaling Republican strength in the congressional elections of 2014. The projections made by different polling organizations turned out to be heavily dependent on the assumptions they made about turnout.<sup>xiv</sup> Different guesses, different results.

The problem here, of course, is that actual turnout is unknown until the election is over. This is likely to be one of the reasons the 2014 polling underestimated Republican strength. Turnout in that midterm election was just 39%--the lowest since World War II. Since Republicans tend to be more committed to voting, a projected normal turnout for a midterm election might well have included more Democrats than actually turned out.

Projecting turnout is especially difficult in primary and caucus states, both because turnout is generally very low, and because different states have different rules about who is allowed to vote in party primaries; some allow independents, some have open registration up till election day, and so forth. For polls, identifying who is actually going to vote can be like finding a needle in a haystack. The first two campaign tests are the caucuses in Iowa and primary in New Hampshire. Historically, only about 20% of those eligible have attended the Iowa caucus. They tend to be more fundamentally religious and conservative voters. The New Hampshire primary has a huge turnout for a primary election, but that is expected to be about 40%. In contrast, one national poll asked if those contacted were likely to participate in the nominating contests and found that almost 70% of the electorate said yes, a tremendous over-report; perhaps one-in-five will before the nominations are decided. One state poll of New Hampshire voters reported that 97% of those they contacted said would vote in that primary, a ridiculous number. Bottom line: there is often a great deal of slippages between a polling organization's modeling of an electorate and what actually happens on Election Day.

## **In Summary**

There are a number of choices to be made in the course of conducting election polling beyond sampling error. We call these "house differences" where different organizations have different ways of doing this type of research. To look for trends, it is probably safest to compare polls done by the same organization at different times, rather than to try to compare polls with different methodologies done at similar times. Given the unique nature of election polling, it is likely that outsiders may look at them with puzzlement and ask "What's going on?"

Some have turned to polling aggregators for help in this task. Polling aggregators are organizations that report the average of polls in a given race, be it national or state. There is indeed some comfort in taking the mean, or statistical average, of polls, as hopefully the outliers are apparent and the sampling error that can lead to higher or lower estimates of what is actually taking place smooths out. There are, however, some cautions that apply here, too. First, polling aggregation works best when there are a large number of polls. Second, some aggregators have minimal standards for the polls they let into their calculations, effectively treating all polls the same, regardless of the underlying level of quality. Third, some of the higher-quality firms, such as Gallup and Pew, have minimal polling schedules in the election, removing a ballast of expertise

from the polling field. Finally, there has been some evidence of “herding,” where polling organizations shape their estimates to agree with the body of already-published polls, which could be compounded through aggregation.

I hope this essay is helpful to our journalistic and other colleagues in understanding some of the sources of variation in election polling. From the inside, those of us conducting election polls see a fair amount of consistency in findings amid the complexity of a science-based-art.

## Why Election Polls May Vary by a Few Percentage Points

- Sampling error
- Length of field period
- Live interviewers vs. IVR
- Type of sample used
- Mode of administration
- Respondent selection
- Likely voter indices
- Question wording & ordering
- Weighting

**Actually, it's  
a wonder  
they are as  
close as  
they are!**

77

---

<sup>i</sup> The reasons for this are laid out in an essay I wrote for the NY Times Sunday Review in the summer of 2015. See [http://www.nytimes.com/2015/06/21/opinion/sunday/whats-the-matter-with-polling.html?\\_r=1](http://www.nytimes.com/2015/06/21/opinion/sunday/whats-the-matter-with-polling.html?_r=1).

<sup>ii</sup> This number is uniformly expressed at a confidence interval of 95%.

<sup>iii</sup> Sampling error is largely based on the sample size, but to a small degree on the percentage division of the sample. A graph showing the relationship between sample sizes and sampling error is appended to this piece.

<sup>iv</sup> Some organizations use a random technique, such as the “last birthday” technique, where the interviewer asks to speak to whoever had the last birthday in the household. There are other techniques of randomization, but the idea is to ensure that everyone has an equal, or at least *known*, chance for inclusion. Other organizations use a systematic technique, such as asking for the youngest male/oldest female at home, that have produced empirically representative samples in the past.

<sup>v</sup> There has been an increasing use of a new sampling frame for probability-based sampling known as **ABS**, for Address-Based Sampling, since the last presidential election. Researchers have access to the master

---

address file of the USPS, through a licensed vendor. A random sample of respondents is pulled from this frame. Names and addresses are run through various data bases to match a phone number with name and address. A possible weakness with this method is in the efficiency of matching numbers to names. This form of sampling is not extensively used in election polling (yet).

<sup>vi</sup> Some unlisted numbers are “findable” when put through various commercial data bases.

<sup>vii</sup> Blumberg SJ, Luke JV. Wireless substitution: Early release of estimates from the National Health Interview Survey, January–June 2014. National Center for Health Statistics. December 2014. Available from: <http://www.cdc.gov/nchs/nhis.htm>.

<sup>viii</sup> This is not meant to indict all surveys without live interviewers. Indeed, such surveys may avoid social desirability biases, such as the tendency to over-report voting intentions that are present to varying degrees in the interviewer – respondent interaction.

<sup>ix</sup> In surveying those reached via a cell phone number in the USA, and assuming the person who answers is an age-eligible adult, most pre-election pollsters proceed to interview the person who answers rather than determining if the cell phone is answered by more than one adult. This is the case because it is thought that few cell phones in the USA among the voting public are shared devices.

<sup>x</sup> Internet Use: <http://www.pewinternet.org/data-trend/internet-use/internet-use-over-time/> Accessed 061615

---

<sup>xi</sup> Report Of *The AAPOR' Task Force On Non-Probability Sampling*, <http://www.aapor.org>  
Reg Baker, J. Michael Brick, Co-Chairs

<sup>xii</sup> There is a handful of Internet-based probability samples, where researchers draw a normal sample and give equipment and Internet service to those who fall in the sample. On-line polls with probability samples include GfK, Gallup, Pew, NORC, USC, and Rand at this writing.

<sup>xiii</sup> Here's an exercise to demonstrate the importance of weighting. On the left is the portion of the electorate, comprised of 25% liberals, 40% moderates and 35% conservatives. In this example liberals split 89-11 for the Democratic candidate, moderates vote Democratic by a margin of 53-47 and 83% of the conservatives vote for the Republican. At these turnout numbers and vote splits, the Republican beats the Democrat by a margin of 51. Holding constant the division of L, M, and C voters, but changing the turnout so that it is comprised of 28% liberals, 40% moderates and 32% conservatives (a change of 3 percentage points in the Ls and Cs) changes the election outcome to a 52-48 Democratic win. In practical terms, one would no more want to weight by ideology than party, but it does show the difference weighing can make if targets change by a few percentage points in some cases.

AGE GROUP	Pct of Electorate	Candidate Choice		Candidate Yield	
		Dem	Rep	Dem	Rep
Liberal	<b>25</b>	89	11	22	3
Moderate	<b>40</b>	53	47	21	19
Conservative	<b>35</b>	17	83	6	29
Total Vote				49	51

AGE GROUP	Pct of Electorate	Candidate Choice		Candidate Yield	
		Dem	Rep	Dem	Rep
Liberal	<b>28</b>	89	11	25	3
Moderate	<b>40</b>	53	47	21	19
Conservative	<b>32</b>	17	83	5	27
Total Vote				52	48

xiv “*Estimating the 2014 National House Vote: What Can Be Learned*” Lydia Saad, Frank Newport, Jeffrey M. Jones and Stephanie Kafka, The Gallup Organization, Presented at the 70<sup>th</sup> Annual Conference of the American Association for Public Opinion Research Conference, May 14-17 Hollywood Florida.

### Relationship between Sample Size and Sampling Error Percentage divisions of 50-50, 70-30 and 90-10 (use blue if unsure)

