

WHY DO THOSE POLLSTERS KEEP TELLING US WHAT TO THINK?

James E. Hanley

Published under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International license. This work may be freely reproduced for non-commercial use under the conditions that it not be changed and attribution is given to the author. 2014.

****DRAFT****

Distrust of Pollsters

Every four years it seems as though public opinion pollsters suddenly appear from the ether like evil spirits determined to invade our thoughts and distort the results to manipulate the presidential election. In fact pollsters are always working, continually polling the public to get the people's perspectives on a wide-range of social and political issues. It's no surprise they're always working - polling is business and businesses need to continually be selling their product.

Polling has grown by leaps and bounds over time, with increasing numbers of polling organizations coming into being and feeding the public numbers about political contests, and many political junkies are desperate to get their daily fix of quantitative election data. But at the same time polls become more ubiquitous and more heavily debated, it seems a growing number of people have come to distrust them, not just from the belief that small samples of people can't tell us what the whole population thinks, but even from a belief that all polls are biased and intended not to reveal the public's beliefs but to shape them in the way the pollster desires. An example of the public's distrust in polling came during the 2012 presidential election. Most polls showed Obama winning (two showed Romney ahead), and Romney supporters interpreted this as evidence of polling bias. Their case was stated with great force and clarity by blogger Tim Stanley.

In the history of presidential elections, has there ever been such an effort by one side to poll their way to victory? While the Republicans have spoken this season about jobs and debt - willing themselves to a moral victory - the Democrats have talked constantly about how well their guy is polling in one or two states. The goal is to create a sense of inevitability, to convince the public to vote for Obama because he's a winner and who wouldn't want to vote for the winner? We've witnessed the evolution of polling from an objective gauge of the public mood to a propaganda tool: partisan and inaccurate.

In the end the polls got the outcome right. In fact, Obama got *more* of the vote than the polls' average predicted. It might seem preposterous that pollsters can decipher what tens of millions of people think just by asking about a thousand of them, but most often they can, although as we'll see it's getting more difficult. The key is to get a sample of the population that reflects the whole country. Imagine a population of clones who all think exactly alike and agree on everything. How many clones would you have to ask a question of before you knew what all of them thought? Only one. Of course people aren't clones, and we don't always agree (otherwise politics would be easy), but we also don't have an infinite variety of opinions on any particular issue. If we have a choice between two candidates, call them Alvarez and Bates, there is a limited variety of opinions: 1) Alvarez, 2) Bates, 3) Neither, 4) No opinion. When the options are that few, you don't need to ask that many people to figure out who the public prefers; you just need to ask a sample that is representative of the population. For example in real-world polling, if you ask only Democrats who they prefer, your answers won't reflect all the people who are likely to vote Republican.

A real world example of how the *number* of people in your sample is not as important as the *representativeness* of the people in your sample comes from the 1936 presidential election. The magazine *Literary Digest* had begun polling the public in order to predict presidential elections in 1920, and had correctly chosen the winner four straight times. In 1936 2.4 million people responded to their surveys, and based on those responses the *Digest* predicted that Republican Alf Landon would defeat Democratic incumbent Franklin Delano Roosevelt. But on election day Landon won only two states, got only 36.5% of the vote to Roosevelt's 60%, and received only 8 of 531 electoral votes. In contrast to the *Literary Digest*, George Gallup, founder of the Gallup Poll, surveyed only 50,000 people, and correctly predicted Roosevelt's victory.

The *Literary Digest* surveyed 50 times as many people as Gallup, so how could he have beaten them? The answer is that you can have large numbers of people who are not representative of the general public. For example you could survey two million homeless people with a question about home ownership, and the result would seemingly show that no Americans owned their own home. Obviously that's ridiculous, but it's similar to what happened in 1936. a. The *Digest* used telephone directories, lists of magazine subscribers, and lists of people who belonged to clubs. But this was during the Great Depression, when most people couldn't afford phones, cars and clubs, so the *Digest* polled people who were wealthier than most Americans, so they undercounted the views of all those people who were not so well off. And it turned out that wealthier people disliked FDR, while less wealthy people liked him.

The irony is that even Gallup used a much larger sample than was necessary: 50,000 people. Today most polls use samples between 800 and 1500 people. The larger size is better, but only up to a point, as I'll show.

How a Small Number of People Can Tell Us About the Whole Country

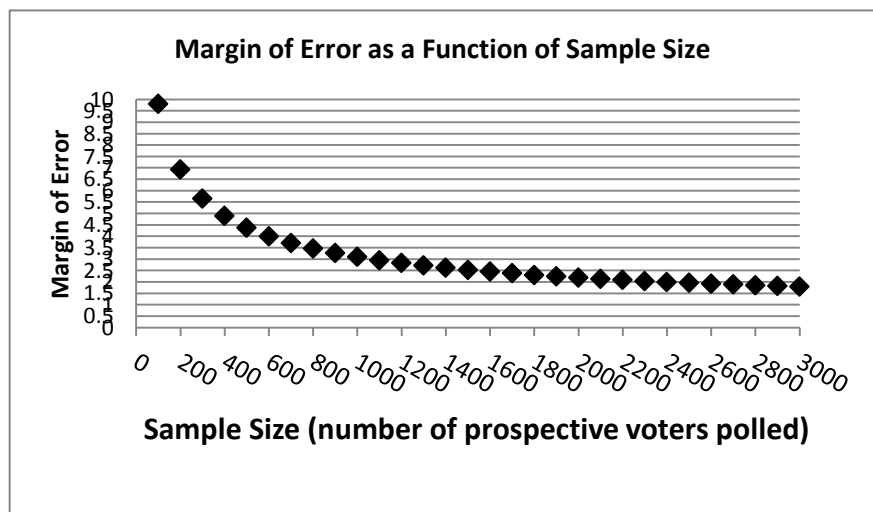
More important than the size of the sample is how the sample is chosen. Imagine a thousand gumballs, blue ones and red ones, in a big jar, and you want to figure out what percentage of

them are red and what percentage are blue. Of course you cannot choose just one gumball and tell how many are blue and how many are red, but at the same time you don't need to count every single gumball. We'll come back to how many you need to count in a moment, but here let's focus on how you choose which gumballs to count. Obviously you can't reach in, look for the color you like, say, blue, grab several of those, and announce "most of them are blue!" But that's exactly what most people do in politics – we surround ourselves with people who are politically similar to us, and then we conclude "Alvarez is going to win the election, because everyone I know is an Alvarez supporter." We also can't just reach in with a cup and scoop some gumballs at the top of the jar, because the gumballs may not be evenly mixed through the jar; all the red gumballs may be on the bottom, and the blue ones on top. But if we make sure all the gumballs are thoroughly mixed up, then our scoop is more likely to get a sample that has a mixture of colors that pretty accurately represents the whole jar.

Mathematicians have figured out how many of those gumballs we would need to count to be pretty sure we know the percentages of red and blue ones. The normal standard of certainty that pollsters aim for is to be (1) 95% certain that (2) the percentages are not off by more than three percentage points in either direction (the margin of error). So if I report that 57% of the gumballs are red, what I really mean is I am 95% certain that the percentage of red gumballs in the jar is between 54% and 60% (57% - 3 percentage points, and 57% + 3 percentage points). So with 1,000 gumballs, how many would I need to count to reach that standard? The mathematical formula says I would need to count 516 of them, over 50%.

But the beauty of sampling is that we don't always have to count half the population. If the number of gumballs increases to 10,000, I have to count 965 gumballs, less than 10%. If the number increases to 100,000, I only have to count 1,056, less than 1%. And if the number of gumballs increases to 1,000,000, the amount I need to count increases only to 1,066. If we increase to 10 million, I only need to count one more, a sample of 1,067, and if we increase that to 100 million—about the number of people who vote in U.S. presidential elections, I only need to add 1 more again, a sample of 1,068. But that only works if my sample of gumballs or people represents the actual diversity within the jar or population of voters.

This can be seen in the following chart, which shows the sample sizes needed to get a particular margin of error, based on the 124 million votes cast in the 2012 U.S. presidential election, and assuming we want to be 95% certain the real views of the public are within that margin of error.



The Pitfalls of Polling

Polling can be done badly, both by accident and purposefully. Even reputable polling organizations sometimes make mistakes - because pollsters are human, too - but sometimes people intentionally use bad techniques to produce a deliberately misleading poll. Organizations that are simply trying to discover the public's beliefs or preferences generally operate honestly, but those that are trying to influence the public, such as campaign organizations and organized interest groups, may dishonestly choose polling techniques that produce a biased result. This section covers the variety of potential errors in polling that can make a poll's results invalid (not truly representative of public opinion).

1. Respondent Knowledge

Respondents (the people who are surveyed) have to have reasonable knowledge and understanding of the subject the pollsters are asking them about. For example, when pollsters asked people whether they supported Obamacare (the Affordable Care Act), a majority said they did not. But when pollsters asked people if they supported the various elements of the policy, without calling it Obamacare, a majority of respondents favored most of them, suggesting that the public as a whole did not really know at that time what was in the law.

Respondent knowledge is equally necessary in campaign polling. If the poll is taken two weeks before a major election asking people whether they support the Democratic nominee or the Republican nominee, we can expect that nearly everyone we might poll will know who and what we're talking about, and will have formed an opinion. In contrast, if we ask people during the early stages of presidential campaigning, more than a year before the presidential election and even months before the first primary, which of the 5 or more candidates they prefer, some of whom they've never heard of, they don't (yet) have the knowledge to give us a considered opinion. As a general rule of thumb, the earlier in the campaign process, the more unreliable polling is, and the later in the process the more reliable it generally is.

2. Question Wording

Writing surveys is something of both an art and a science. Good survey companies have learned from both experience and experiments how to write questions that respondents can understand and that don't lead them to a particular answer (either accidentally or purposefully). But pollsters are humans, so sometimes they make mistakes. And some are unethical and purposefully try to trick respondents into giving a particular answer.

Confusing Wording

A survey done in 1992 demonstrates how even a good polling organization sometimes makes mistakes and writes their questions badly. Asking respondents about the Holocaust, the Roper polling firm asked,

“Does it seem possible, or does it seem impossible to you that the Nazi extermination of the Jews never happened?”¹

If you find that question confusing, you are not alone. It created an uproar, because almost a quarter of the people polled replied “possible it never happened.” But the question didn’t show what people really believed - it just confused them with its awkward phrasing of the question (poll questions should never have a double negative). A follow-up poll with a more clearly written question showed that only 1% of the public doubted that the Holocaust happened. The badly written question was not done purposefully, but shows that even reputable organizations sometimes make mistakes.

Unexpected Word Influence

In 2010, polling revealed that 59% of people favored allowing homosexuals to serve in the military, but 70% favored allowing gay men and lesbians to serve in the military. The question was identically phrased in both polls, except for the particular word used to describe gays and lesbians. While we know that such word choices can matter, we cannot predict when they will.

Leading Questions

A leading question is one that biases the respondent toward a particular answer. Reputable polling firms carefully review their surveys to try to avoid leading questions, because the answers are unreliable if what you want to know is what people really believe. But un reputable pollsters will write them intentionally so they can falsely report the public’s beliefs. This dishonest type of polls is called a “push poll,” because it is trying to push people to answer in certain ways.

“Do you believe Republican leadership has solutions for anything?”²

This question is leading because it’s essentially a sneer. It suggests to the respondent that in fact Republicans don’t have any solutions for any conceivable political issues. Good poll questions do not imply what the “correct” answer is.

“Harold Ford voted against renewing the Patriot Act, which treats terrorists like terrorists, but Bob Corker supported renewing it. Based on what you now know, do you intend to vote for Harold Ford or for Bob Corker?”³

This question is leading because it assumes that a particular political opinion is correct (that we should treat terrorists like terrorists) and identifies which candidate opposes and which supports that opinion before asking the respondent who they support. The purpose is to lead people to the “correct” answer.

As a general rule, if the questions of a poll seem to suggest how you’re supposed to answer them the poll is probably a push poll and should be avoided and not trusted, even if seems to support your political views.

3. Biased Samples

As we've already seen, the people who take the survey have to be representative of the general public. Of course no individual person can be representative of such a diverse group of people, but all the respondents, as a group, need to be reflective of the whole public. A non-representative group is what researchers call a *biased sample*, and avoiding a biased sample can be difficult, for several reasons.

Low Response Rates

According to Pew Research, one of the premier research firms, the response rate—the percentage of people called who actually agree to take and complete a survey fell from 36% in 1997 to 9% in 2012.⁴ Part of this decline probably has to do with a decline in trust in both pollsters and in government itself, but a lot of it has to do with the shift from landlines to cell phones - by 2015 over 60% of Americans could only be reached by cell phone.⁵ Most landline phones had no caller identification, so people had to answer the phone to see if they wanted to take the call, and once people are actually speaking to a real person they find it harder to hang up. Cell phone users frequently ignore any numbers they don't recognize, making it harder to reach them. Additionally, the federal Telephone Consumer Protection Act of 1991 prohibits using automatic dialing equipment to call cell phone, so calls to cells must be dialed by hand, increasing the costs for pollsters.

Low response rates don't necessarily mean a biased sample. They can just mean the polling organization has to make many more calls to get the sample size they need. For example, if you need a sample of 1,000 people and the response rate is 36%, you'll need to call almost 2,800 people. But to get 1,000 respondents with a response rate of 9% you would have to call over 11,100 people.

And a good poll has to make the effort to get people with cell phones to answer and take their surveys, because people who use only or mostly cell phones are different than people who use only or mostly landlines. According to one researcher they tend to be younger and more Latino, and as a group are 16 percentage points more likely to be Democratic.⁶ A poll that didn't make the extra effort to get a sufficient number of cell phone users would overestimate support for Republican candidates and underestimate support for Democrats. For example, the same researcher notes that in the 2012 Indiana Senate Race, the Republican candidate's own polls showed him with a two-point lead, while the researcher's poll showed him trailing by seven points. In the end he lost by six points.⁷ The likely reason for the candidate's flawed polls is that his polling team was only calling landlines, probably because of the lower cost.

But high quality polling firms don't just dial lots of numbers until they get someone to answer. They will call each number multiple times, at different times, over a period of several days. If, for example, they only called people between 5 and 9 p.m. on weekdays, the

polls would exclude people who work evenings. So to make sure those people are included, too, if a number is selected for calling but nobody answers, the next call will be made at a different time of day, and this will be repeated at different times for several days. Often 5-6 attempts are made before giving up on that number. This process is described in an article about the firm Public Policy Polling.

On the Friday after the Democratic convention, Tom Jensen tried to reach out and touch 10,000 Ohioans. He wanted to ask them, among other questions, whom they planned to vote for in November: Barack Obama or Mitt Romney? This sort of thing is easier—and harder—than you might think. As the director of Public Policy Polling, Jensen has at his disposal 1,008 phone lines hooked up to IVR (interactive voice response) software that enables PPP to make 400,000 automated calls a day. All Jensen needs to do is feed the 10,000 phone numbers into a computer, record the series of questions he wants to ask, press a few buttons, and voilà: He has a poll in the field. That's the easy part. The hard part starts with getting people to answer the phone. Beginning that Friday night around six and then five more times over the course of the next two days—in the mornings, afternoons, and evenings—PPP called those 10,000 Ohioans; by Sunday night at eight, only 1,072 of them had been reached. Still, for Jensen's purposes, that was sufficient, and he got to work assembling his poll.⁸

If the people who don't answer are politically similar to those who do answer, the only problem is that it takes more time to find enough respondents - the results won't differ. But if the people who don't answer tend to be politically different than those who do answer, then our poll results today will be less accurate.

Making Statistical Adjustments for Non-representative Samples

One of the ways to solve the problem of a non-representative sample is to make statistical adjustments. This method is controversial among the public, to whom it seems like illegitimate manipulation of the results to achieve a pre-determined result, but among pollsters it is recognized as a legitimate approach, although there is debate about just how to do such statistical adjustments. This is analogous to a group of doctors examining a sick patient - they all agree the patient needs treatment, but they may not agree on just what that treatment should be.

For example, if a survey sample has 60% women and 40% men, we know that doesn't represent their proportions in either the general population or the voting population. If men and women are the same in their political views or who they are likely to vote for this wouldn't matter, but we know that statistically men and women do differ. Men are more likely than women to vote Republican, women are more likely to be pro-choice, men are more likely to support the death penalty, and so on. In our example, if we did not do statistical adjustment to account for the lower proportion of men in our sample we would

underestimate the Republican vote, overestimate support for abortion rights, and underestimate support for the death penalty.

Pollsters also may purposely try to sample proportional amounts of particular groups to ensure they get a representative sample. For example, Latinos are about 17% of the U.S. population, but only about 9% of voters. If we are just surveying Americans about political attitudes, we would want to match their proportion in the population, but if we are trying to predict the outcome of an election, we don't want 17% of our sample to be Latinos, only about 9%. But when they cannot get a sample that proportionately represents the various groups, statistical analysis is the tool to increase accuracy.

One controversy that has been common in recent election cycles is alleged oversampling of Democrats and statistical adjustments to the number of Democrats polled. Both in 2012 and 2016 supporters of the Republican candidates claimed bias and fraud in polling because many polls either sampled more Democrats than Republicans or statistically adjusted their results to give more weight to their Democratic respondents. If the American public was split evenly between Democrats and Republicans, this would be a fair complaint. But in fact more Americans identify as Democrats than Republicans, so an even number of each in the sample would improperly undercount Democrats. Because the pollsters are trying to accurately predict the outcome of an election, they want their samples to mirror the electorate as closely as possible. If some day the American public shifts so that Republicans outnumber Democrats, then pollsters will shift to ensuring they have appropriately more Republicans in their samples.

Internet Polling

In response to the increasing difficulty of getting people to talk to pollsters on the phone, some polling organizations are experimenting with using the internet. The organization YouGov, for example, created a base of 1 million people who agreed to participate, for payment, in internet surveys when selected, and each time they conduct a survey they select around 1,000 of them. The University of Southern California and the Los Angeles Times Newspaper collaborated to do something similar. These approaches are experimental and it is too soon to tell how well internet-based polling will work, but polling is a profitable business, so polling organizations have a financial incentive to keep experimenting until they find techniques that work well.

Self-Selected Samples

Television stations, newspapers, and websites often have polls where they invite people to call in or click on a button to participate. These polls are not to be trusted because the samples are wholly self-selected and inevitably produce biased samples and unreliable results. Often people can participate as many times as they want to. Some of these polls are honest enough to include a note that their poll is not scientific - a warning that should be

taken seriously! - but not all bother to do so. Regardless, you should view all polls like these as junk and not to be taken seriously.

How Citizens Should Evaluate Polls

Although good polling has become difficult enough that it can be said we are in “polling’s dark age,”⁹ it is still largely an effective way to gauge public opinion, and citizens who understand how to evaluate polls can make better sense of the political landscape than those who do not. Fortunately, understanding polls is not terribly hard. We just have to pay attention to who’s doing the polling and how they’re doing it.

1. Pay Attention to Who Is Doing the Polling

Polling should try to accurately gauge the public’s beliefs, but can also be done to mislead and perhaps shape their beliefs. Poll results from ideologically motivated organizations or any political candidate’s campaign organization should be viewed skeptically. Campaign organizations are notorious for overstating public support for their candidate, especially when that candidate is losing. They may do this because they use bad polling methods, or they may actively lie, or they may just make errors through wishful thinking. This occurred with the Mitt Romney presidential campaign in 2012, where the campaign’s pollsters showed Romney winning. It appears that they believed the major polls oversampled Democrats. Other polling showed Republicans as more motivated to vote than Democrats, so the campaign pollsters oversampled Republicans to get what they thought was a sample more representative of actual voters - their wishful thinking misled them.¹⁰

The major polling firms do a much better job. They are trying to make money by selling their results to news organizations, and their credibility is key. Their reputation is their real product. But even they make mistakes. Every organization, from a multi-national corporation to a neighborhood church can decline in quality as management changes hands, and polling organizations are no different. Two notable polling organizations have been criticized for poor performance in recent years, one of them following a change of ownership that caused the poll’s founder to leave the organization. But what declines often rises again, and the next campaign could see both those organizations having corrected their mistakes, because they have a financial incentive to do so.

2. Look at the Sample Size and How the Sample is Selected

Look at the sample size and how the sample was selected. If the sample size is below 800 or so that means the margin of error is large (see below). That doesn’t mean the poll is *wrong*, just that it is unlikely to be precise. The best polls normally say that their sample was selected randomly. If the poll says nothing about how the sample was selected, be wary of it.

3. Look at whether the Poll is Sampling the General Public, Registered Voters or Likely Voters

A poll about political beliefs does not need to limit itself to registered voters, but if the poll is trying to predict an electoral outcome then it should not sample people who are not eligible to vote. So at a minimum an electoral poll should say that it sampled registered voters. But not all registered voters actually do vote, so pollsters often ask people if they are likely to vote. This focuses the poll on likely voters, giving a better prediction of the electoral outcome.

4. Pay Attention to the Margin of Error

The margin of error gives us the range by which our poll results might vary. This means if candidate Adams leads candidate Baker 52% - 48%, and our margin of error is ± 3 percentage points, we're pretty sure that if we could poll all voters instead of just a sample of them the real range would be somewhere between Adams leading 55%-45% and Baker leading 51%-49%. But the numbers closer to our poll result are more likely than numbers farther away, so it is more likely that Baker really has the support of 53% of the voting public than 55% or 49%. So in our example of Adams leading 52%-48%, it is possible that he could be tied or even losing, but it is more likely that he is in fact leading.

As we saw above, the margin of error is determined by the sample size. Major polls usually select a sample size that will give them a margin of error of ± 3 percentage points. There is nothing inherently wrong with using a smaller sample and settling for a larger margin of error. Sometimes we don't need as much precision. For example, early in the election process, when a party has multiple candidates vying for the nomination, we know that their levels of support will change over time so it's too early to predict a winner. So we might just be polling people to see which candidates have a substantial level of support and which have very little support. In that case a larger margin of error is sufficient, and sampling more people to get more precision would cost more without giving us the benefit of more meaningful information.

Because of the margin of error we cannot always say when a candidate's support is rising or falling. If we have a margin of error of ± 3 percentage points, and our first poll shows a candidate with 46% support while our second poll shows them with 48% support, that doesn't mean the candidate has actually gained more support from the public. They might have, but because the difference between the numbers is within the margin of error, we can't know that. It's even possible the candidate has slightly lost support, but our measurements weren't precise enough to capture the change.

5. For Presidential Elections, Pay Attention to State Polls

While national polls are usually fairly accurate, the presidential election is not based on the popular vote but on the electoral college. When the popular vote is not close, the national polls can be expected to predict the winner. But in a close presidential race, what really matters is which states each candidate wins, in order to secure a majority of electoral votes. In that case the interested observer should look closely at the polls in crucial states.

6. Look at an Average of Polls

Any poll can be off-base sometimes. Generally we're 95% sure that the real beliefs of the public are within a margin of error of ± 3 percentage points, but that means there's a 5% chance we're wrong. If so, the poll is much more likely to be only slightly wrong rather than very wrong. That's one of the reasons why different polls often produce different results. So rather than look at a single poll we should look at multiple polls and average them. The collection of polls is more likely to be accurate than any single poll, because those that are wrong in one direction are likely to be offset by those that are wrong in another direction.

What many people do is focus on the polls that show their preferred candidate winning, and assume - or maybe just hope desperately - that those are the right ones. But they have no real reason to believe those polls are better than the others (assuming all of them are reputable organizations). The objective political analyst, as opposed to the mere political advocate, will look at the average of polls rather than letting themselves be seduced by the ones that show them what they want to see.

Additionally, in a close election where most of the polls show the candidates' support as within the margin of error, look to see if one candidate is leading in a large majority of polls. If about half the polls show one candidate in the lead and about half show the other candidate in the lead, then the election is a true tossup. But even in a close election, if the large majority of polls show the same candidate in the lead, that candidate is almost certainly the true leader, at least at that moment.

Prediction Markets

A more recent innovation for forecasting electoral outcomes is prediction markets. In these markets people bet money on electoral outcomes. A common practice is to have every payoff be either \$0 (if you lose) or \$1 (if you win). For example, if I think candidate Adams will win I can buy a share at the going price. Let's say that when I buy that price is 30 cents. If Adams wins, I'll get \$1 for that share, meaning a gain of 70 cents. If Adams loses, I get \$0 and lose my 30 cents.

What that 30 cent share means is that the collective wisdom of those buying shares in the market think Adams has a 30% chance of winning. If I think her chance is greater, I'll be willing to spend more than 30 cents to get a share. If someone else thinks her chance is less, they'll be happy to sell me a share. The market price for shares at any given moment is expected to accurately reflect the real probability of the candidate winning.

If that sounds improbable, experience shows that prediction markets are often more accurate than polls.¹¹ The essential idea is that if people put their money where their mouth is, they won't risk it unless they have enough knowledge to do so. Economist Alex Tabarrok calls betting "a tax on bullshit."

Overall, I am for betting because I am against bullshit. Bullshit is polluting our discourse and drowning the facts. A bet costs the bullshitter more than the non-bullshitter so the willingness to bet signals honest belief. A bet is a tax on bullshit; and it is a just tax, tribute paid by the bullshitters to those with genuine knowledge.¹²

Prediction markets are not perfect, of course. Just as people sometimes over or under value companies in the stock market they sometimes over or under value candidates or issues. In 2016 the prediction markets underestimated Donald Trump in the presidential primary. And in the same year, when Britain voted to leave the European Union, the prediction market showed a false expectation that voters would choose to stay.¹³ This does not mean prediction markets are useless. What it means is that careful political observers should pay attention to both polls and prediction markets as they both provide relevant information. The more closely they agree, the more certain we can be in the outcome, and when they disagree we should be more cautious about assuming we can predict what will happen.

Conclusion

Polls are meaningful measures of public opinion, and don't deserve the degree of distrust many citizens have for them. But polling is tricky to do well, and it is easy to do it dishonestly. Citizens should pay attention to how polls are conducted, how questions are written, the sample size and the method of selecting the sample, and the margin of error. They should also look at a large number of reputable polls and average their results rather than rely on the accuracy of any particular poll. Further, they should look at prediction markets and observe which outcome the tax on bullshit favors.

¹ Roper Organization. 1992. Poll for the America Jewish Committee.

² Ed Schultz Show. 2010.

³ Robo Push Poll in the 2006 Tennessee Senate Race.

⁴ Pew Research Center. 2012. "Assessing the Representativeness of Public Opinion Surveys." <http://www.people-press.org/files/legacy-pdf/Assessing%20the%20Representativeness%20of%20Public%20Opinion%20Surveys.pdf>.

⁵ Marketing Research Association. 2015. "The Telephone Consumer Protection Act (TCPA) and Calling Cell Phones." <http://www.marketingresearch.org/legal-article/tcpa-restrictions-using-autodialers-call-cell-phones>.

⁶ Stryker, Brian. 2014. "Can We Trust Polls Without Cellphones Anymore?" *The Huffington Post*. http://www.huffingtonpost.com/brian-stryker/can-we-trust-polls-withou_b_4880127.html.

⁷ Ibid.

⁸ Zengerle, Jason. 2012. "The. Polls. Have. Stopped. Making. Any. Sense." *New York Magazine*, Sept. 30. <http://nymag.com/news/features/election-polls-2012-10/>. Accessed October 30, 2016.

⁹ Hawkins, Beth. 2012. "What's the Problem with Polls These Days? Experts Take a Stab at Answering." *Minnpost.com*. <https://www.minnpost.com/politics-policy/2012/10/whats-problem-polls-these-days-experts-take-stab-answering>.

-
- ¹⁰ Crawford, Jan. 2012. "Adviser: Romney 'Shellshocked' by Loss." CBS News. <http://www.cbsnews.com/news/adviser-romney-shellshocked-by-loss/>. Accessed April 4, 2016.
- ¹¹ Posner, Richard. 2008. "Prediction Markets and the Election." *The Becker-Posner Blog*. Sept. 14. <http://uchicagolaw.typepad.com/beckerposner/2008/09/prediction-markets-and-the-election--posner.html>. Accessed October 30, 2016.
- ¹² Tabarrok, Alex. 2012. "A Bet is a Tax on Bullshit." *Marginal Revolution*. Nov. 2. <http://marginalrevolution.com/marginalrevolution/2012/11/a-bet-is-a-tax-on-bullshit.html>. Accessed October 30, 2012.
- ¹³ Rothschild, David. 2016. "Brexit: What Happened?" *PredictWise*. June 24. <http://www.economist.com/blogs/graphicdetail/2016/06/polls-versus-prediction-markets><http://www.economist.com/blogs/graphicdetail/2016/06/polls-versus-prediction-markets>.