Election Polls are 95% Confident but only 60% Accurate

Aditya Kotak
University of California, Berkeley

and

Don A. Moore
Haas School of Business
University of California, Berkeley

In press at *Behavioral Science and Policy*

Abstract

Public election polls are more confident than accurate. We find that reported 95% confidence intervals are only 60% accurate, even for polls conducted in the week prior to the election. These polls would in fact need margins of error twice their reported width in order to be 95% accurate. When lay people read about poll results they tend to overestimate polls' accuracy, even when they get historical data on polls' overconfidence. These results illustrate polls' vulnerability to overconfidence and limitations to the lay public's understanding of these shortcomings. We conclude by providing commentary on the challenges of calibrating confidence in an uncertain future, especially when communicating with the general public.

Author Notes

This project's pre-analysis plans, data, and source code are available: <https://osf.io/65za7/>.

Thanks to Amelia Dev for her support for this project, and to Gabriel Lenz and David Broockman for valuable feedback.

Introduction

In 2016, the world was stunned by a couple of surprising election outcomes. On the 23rd of June, 52% of voters in the United Kingdom elected to leave the European Union. Many eligible voters who supported remaining in the EU did not bother to vote that day, possibly because most credible polls forecast a likely victory for the “remain” side (YouGov, 2016). Voter turnout was lowest in those areas, like London, Scotland, and Northern Ireland, that most strongly favored remaining in the EU (BBC News, 2016). Then, in November, Donald Trump beat Hillary Clinton in a closely contested U.S. Presidential election. On the eve of the election, poll aggregator Nate Silver’s FiveThirtyEight.com web site gave Clinton a 70% chance of winning (Silver, 2016a). Many Democrats who did not vote reportedly believed their votes did not matter (Harvey, 2016).

Polls are important, not just because they influence voter turnout. The Commission on Presidential Debates only allows the participation of candidates achieving more than 15% support in polls (The Commission on Presidential Debates, 2020). Candidates leading in the polls attract more support, including financial contributions (McMinn & Hurt, 2020). And polls are powerful drivers of press attention and its notorious “horse race” coverage (Westwood et al., 2020). Pre-election polls have increased in number and frequency since their introduction by George Gallup in 1936 (Hillygus, 2011; Jennings & Wlezien, 2018). Polls garner attention from the public, the press, and politicians eager to divine the will of the electorate (Ford, Wlezein, Pickup, & Jennings, 2017). Given polls’ power in democracies, all of us ought to care about their accuracy.

This project set out to compare pollsters’ confidence with their accuracy. Polls are reported with margins of error that quantify the uncertainty around a poll’s result. When a margin of error is reported with a 95% confidence level, this provides a 95% confidence interval. This confidence interval makes the claim that there is a 95% likelihood this interval truly captures the population’s vote proportion. It is common practice for polls to sample 800 likely voters and

estimate a 95% confidence interval around the poll's result, yielding a margin of error of plus or minus 3.5 percentage points. So, for example, a poll finding that 50% of 800 likely voters favor a particular candidate in an upcoming election, then a 3.5% margin of error would imply a confidence interval predicting 95% chance that the candidate will receive between 46.5% and 53.5% of the vote in the actual election.

In this paper, we investigate polls' accuracy by testing how often their confidence intervals include the actual election outcome. We also build on this result by considering how much confidence the public has in election polls and how reporting methodologies could help improve public interpretation of polls.

Sources of Error

Margins of error that specify a 95% confidence interval usefully quantify uncertainty due to sampling error. That is, because polls include only a subsample of the population, it is possible that random sampling yields a biased estimate. If pollsters just happened to reach more Clinton supporters ahead of the 2016 election, then if 60% of the poll's respondents favored Clinton we cannot be sure that this 60% predicts Clinton's actual support on election day. How wide should our confidence interval around the 60% estimate be? Statistics helpfully quantify confidence interval width, assuming the error in our estimate arises from failures of random sampling due to chance.

However, the literature on total survey error documents at least five additional sources of error, more difficult to quantify (Biemer, 2010; Groves & Lyberg, 2010). Specification error is the result of a mismatch between the survey question and the forecast target. For example, the question, "Which candidate is better qualified?" is poorly specified to predict voting preferences, since not everyone votes for the candidate they believe to be best qualified. Frame error describes the discrepancy between the sample and the population; respondents randomly sampled at a Clinton rally would poorly predict the broader electorate. Nonresponse error arises

from nonrandom nonresponse; for example, Trump supporters' suspicion of the mainstream media may have made them reluctant to respond to polls. Measurement error describes bias introduced by the method of measurement; the race of an interviewer might, for example, affect respondents' reported racial attitudes. Finally, data-processing errors can occur in the cleaning or analysis of poll results.

Even pollsters who acknowledge the existence of these sources of error may have difficulty using that knowledge to adjust their confidence intervals. The dilemma is captured by data-processing errors. Every researcher tries to minimize them, but knows that they occur. How much have data-processing errors influenced the results of a particular poll? Answering that question depends on the pollster knowing what errors she has made; but if she knew that, then she would have corrected them. An alternative approach might attempt to quantify base rates of data-processing errors. Here again, it can be difficult to compute global error rates without a useful estimate of any local error.

The overwhelming response to the difficulty quantifying these aspects of total survey error is to ignore them. That is, pollsters do their best to minimize these different sources of error and then pretend that they have succeeded. Reported confidence intervals include known sources of error due to sampling, but omit others (Voss et al., 1995). The consequence is that reported confidence intervals are likely to be too small (Goot, 2021). They will be overly precise and 95% confidence intervals will include the actual election outcome less than 95% of the time. Our first study asks how well polls actually predict election outcomes. Our second study reports what lay people believe about polls' accuracy.

Study 1: Polls and Election Forecasts

While polls are not technically forecasts, they are routinely treated as such (Campbell & Lewis-Beck, 2008). If polls are more confident than they are accurate, they would be in good company. Overprecision in judgment is one of the most pervasive biases; it afflicts most human

judgments, including forecasts. Research suggests that forecasters routinely behave as if they are too sure that they know what will happen (Kahneman & Tversky, 1973; Moore et al., 2017). They are most likely to be overconfident when they underestimate their vulnerability to error and neglect to consider all the ways in which they could be wrong (Fischhoff, Slovic, & Lichtenstein, 1977).

Tetlock (2005) distinguishes approaches to forecasting by their totem animals, drawing on the ancient Greek aphorism, “The fox knows many things, but the hedgehog knows one big thing” (Berlin, 1953). Those experts who bring a strong ideological orientation and neatly fit the messy details of reality’s complexity to their overriding narratives and clear conclusions are hedgehogs. Foxes, by contrast, are less likely to see universal laws and coherent ideologies; they are open to revising their views and accepting the possibility that they might be wrong. Perhaps it should be no surprise, then, that foxes consistently make more accurate forecasts than do hedgehogs (Tetlock & Gardner, 2015). After all, perhaps the most useful strategy for debiasing overconfidence is to consider the possibility that you might be wrong (Lord et al., 1984).

Research examining poll accuracy has primarily assessed the proximity of poll results to the vote-shares candidates receive in the election (Ford et al., 2017; Jones Jr, 2008). We seek to contribute to prior research by comparing accuracy against confidence, thereby testing calibration. We investigate the calibration of polls and their margins of error because they afford a benchmark against which we can assess accuracy. Our second study then invites laypeople to report their faith in the accuracy of polls and compares those reports with Study 1’s results.

Method

We preregistered a plan to investigate whether election polls’ confidence levels reflect their actual accuracy. We were able to obtain data from RealClearPolitics.com on 1,931 polls in 4 election cycles: Democratic presidential primaries in Iowa and New Hampshire from 2008,

2016, and 2020 as well as Republican primaries in the same states from 2012, 2016, and 2020.

We also obtained polls conducted ahead of the general elections of [2008](#), [2012](#), [2016](#), and [2020](#). Prior to 2008 the data did not consistently include sample sizes. Because some polls asked about several candidates, these 1931 polls produce 6654 vote share estimates with accompanying margins of error. Our preregistration, data, and code are available at <https://osf.io/65za7/>.

We also included primary elections, but only Iowa and New Hampshire, the first two primaries, because beyond that the race gets more complicated. For instance, candidates may drop out of the race between the poll and the primary election making it difficult to assess the accuracy of forecast vote shares for the absent candidate, and changing the competitive landscape for remaining candidates. We also could not use national primary polls as there is no national primary election, so we would have no result to compare these polls against. We also exclude polls that provide only a point estimate without a margin of error. These constraints left the following primaries:

- [2020 Iowa Democratic Caucus](#)
- [2016 Iowa Republican Caucus](#)
- [2016 Iowa Democratic Caucus](#)
- [2020 New Hampshire Democratic Primary](#)
- [2016 New Hampshire Republican Primary](#)
- [2016 New Hampshire Democratic Primary](#)
- [2012 Iowa Republican Caucus](#)
- [2012 New Hampshire Republican Primary](#)
- [2008 Iowa Democratic Caucus](#)
- [2008 New Hampshire Democratic Primary](#)

In total, we analyze data for 14 election cycles' worth of polls which yielded 1931 polls and 6654 forecasts. The modal poll reported the 95% margin of error of 1.96 standard deviations around the mean, assuming a normal distribution in sampling error. For those polls that did not report margins of error, we used the poll's result and sample size to compute a 95% margin of error.

Results

Our preregistered analysis plan called for three analyses, all of which test the accuracy of polling results as a function of time to the election: (1) regressing the probability of the forecast likelihood of the election outcome on days between the poll and the election; (2) regressing the deviations of the poll outcome from the election result on days between the poll and the election; and (3) a logistic regression testing the rate at which polls “hit” (by getting the election result inside the margin of error) as a function of days to the election. The analyses we present here are more informative than, but consistent with, these preregistered analyses. Because we believe that the preregistered analyses are less useful, we relegate them to a supplementary results file, posted online at <https://osf.io/keswd/>.

The size of a margin of error is a function of a confidence level, the sample variance, and the sample size. A poll “hit” if the poll’s 95% confidence interval included the election’s actual result. Our primary analysis examined hit rate as a function of temporal distance between the poll and the election.

We grouped the polls in seven day intervals and averaged the hit rate of polls in each interval. A linear regression weighted by the number of polls for the given week provides an estimate of the actual accuracy of polls as a function of time until the election. Figure 2 shows that hit rates average around 60% in the week prior to the election. A year prior to the election, average hit rates decline to around 40%.

Figure 2. Hit rates as a function of weeks until the election.

Since most polls ask “If the election were held today...” it might be no surprise that accuracy goes down as the distance in time between the poll and the election goes up. However, even just a few days before the election, 95% confidence intervals only capture the actual vote share only 60% of the time.

We tested how much wider the confidence intervals should be in order to be 95% accurate. This analysis sampled point estimates classified as “misses” from each weekly group and calculated how much wider the interval would need to be for each week such that 95% of the point estimates would include the true election outcome. Figure 3 visualizes our findings. A week before the election, reported margins of errors would have to be 2 times wider in order to be 95% accurate. A year before the election, margins of error would have to expand by more than a factor of 3 to be 95% accurate.

[Insert Fig 3 about here]

Figure 3. Adjustment required to 95% confidence intervals for polls to attain 95% hit rate.

Finally, we also compared election cycles to look for trends over time. We conducted the same analysis as shown in Figure 2 but segmented by year. Figure 4 does not support the idea that polls have become less accurate over time.

Figure 3. Accuracy of polls, by election year, as a function of distance from the election. Shaded region shows the 95% confidence interval around the 2020 results.

Study 2: What Do Laypeople Think About Polling Accuracy?

Given the low rates of accuracy, we wanted to know how skeptical laypeople are when they read reports of political polls. Do they understand polls' actual accuracy for predicting election outcomes? We preregistered a plan to test seven different poll reporting styles:

<https://osf.io/9qhmf>

Method

We planned a sample size of 200, recruited from the Cloud Research list of approved participants from Amazon Mechanical Turk (Litman et al., 2017). We restricted our sample to residents of the United States, seeking a population that roughly matches the demographics of the US (Moss & Litman, 2020). In order to obtain 200 valid responses given our preregistered exclusion criteria, we opened our survey to 230 and wound up with 217 complete responses.

Data, materials, and code are available: <https://osf.io/5wmqe/>

The survey presented each respondent with seven different reporting styles for the same poll result, in the following order:

1. Point estimate (Candidate A polled at 49%)
2. Point estimate with margin of error (Candidate A polled at 49% +/- 3 pts)
3. Point estimate with margin of error and 95% CI (95% confidence that Candidate A's true vote share will be 49% +/- 3 pts)
4. Interval (Candidate A polled at 46% - 52%)
5. Interval with 95% CI (95% confidence that Candidate A's true vote share will be between 46% and 52%).
6. Point estimate with margin of error, 95% CI, and empirical accuracy (60% at one day, 55% at three months, and 35% at one year)
7. Interval with 95% CI and empirical accuracy (60% at one day, 55% at three months, and 35% at one year)

For each of the seven reporting styles, we asked respondents to indicate, on a 0-100% scale, how confident they were that the election outcome would be consistent with the poll's forecast.

We also sought to test how the time between a poll and the election affected respondents' confidence and so we randomly assigned participants to one of three time horizons:

- A) one year prior to the election
- B) three months prior to the election
- C) one day prior to the election

We predicted that confidence in polls would exceed their historical accuracy.

Results

As predicted, average confidence ($M = 59.9\%$) exceeds average accuracy ($M = 49.7\%$), $t(216) = 6.81, p < 10^{-11}$. But confidence in a poll's predictive accuracy varies by time horizon and reporting style. A 3 (time horizon) X 7 (poll reporting style) mixed ANOVA with repeated measures on the second factor reveals a main effect of time horizon, $F(2, 1491) = 30.1, p < 10^{-13}$, and a main effect for reporting style, $F(6, 1491) = 7.6, p < 10^{-7}$. Their interaction is not significant, $F(12, 1491) = 1.38, p = .17$.¹ See Figure 5.

Figure 5. Confidence in poll results, by time horizon (designated by different line patterns) and reporting style (x-axis). Error bars show standard errors.

While we found that reported confidence exceeds historical accuracy across all reporting styles, this was not universally true for all time horizons. Style 2 represents one of the most common: a point estimate, plus or minus a margin of error. For reporting style 2, confidence exceeds accuracy only for the one-year time horizon, $t(74) = 8.5, p < 10^{-11}$, and not for three months, $t(71) = 1.1, p = .29$, or one day, $t(69) = 2.39, p = .02$. Figure 5 underscores both people's excessive faith in polls' predictive accuracy and the challenge of correcting it. Providing more information about polls' poor record of accuracy in reporting styles 6 and 7 failed to bring confidence into line with historical accuracy.

General Discussion

We find that pre-election polls' 95% confidence intervals include the actual election result substantially less than 95% of the time. In order for their confidence intervals to be calibrated with their hit rates, those confidence intervals would have to at least double in size.

¹ We erroneously only preregistered the one-way ANOVA examining the within-subjects effect of reporting style. This within-subjects main effect is significant, but the omnibus ANOVA we report is more useful.

Without this corrective, we find that variation in the styles by which poll results are reported fail to correct superficial assumptions about their accuracy. That is, even when informed about polls' low historical accuracy, people continue to have excessive faith in the current poll's predictive accuracy. They report believing that polls' margins of error are more likely to include the truth than historical averages suggest, although we can take small comfort from the fact that reported confidence is also below 95%.

Why are polls so inaccurate? Poll results might deviate from election results because of errors due to sampling, frame, nonresponse, measurement, specification, and news (Shirani-Mehr et al., 2018). The statistics we use to build confidence intervals have an easier time accounting for sampling error than for other sources of error arising from systematic differences between the poll's sample and the voters. Because statistical models have trouble quantifying these differences, computed confidence intervals will appear overconfident. That is, computed confidence intervals imply excessive confidence in the accuracy of the poll.

History includes some surprising election outcomes. The 2016 election of Donald Trump to the U.S. presidency is just one example. That outcome prompted speculation about declining accuracy in polling (Cohn, 2017; Mercer et al., 2016). It is possible that the 2016 election reflected a novel phenomenon, such as an increased reluctance by right-wing voters to respond to polls--a reluctance born of their more general suspicion of media organizations (Russonello, 2020). If that were the case, polling errors should be increasing. However, we do not find that hit rates were lower for more recent election cycles per Figure 4 (see also Jennings & Wlezien, 2018).

The more parsimonious explanation for the 2016 result is that it represented an aberration, and that both polls and elections will always include noise. Although many observers criticized Nate Silver and his FiveThirtyEight.com election forecasting site for giving Hillary Clinton a 71% chance of prevailing on the eve of the 2016 election, Silver (2016b) sensibly defended himself by noting that events with a 29% probability of occurring do happen.

Nevertheless, Trump's surprisingly strong showing four years later in the 2020 election underscored concerns that recent polling errors may be the result of a larger social trend that ought to drive the interpretation of polls going forward. Going into the 2020 presidential election, national polls favored Joe Biden by 8 percentage points (Silver, 2020). In fact, he only won by 4.5% (Russonello, 2020). Were Trump's surprisingly strong performances in 2016 and 2020 due to chance? Or might they be attributable to a "Trump effect" (Isakov & Kuriwaki, 2020)? More importantly, do they reflect long-term trends with implications for future polls? These questions highlight the challenge of adjusting polls to account for their shortcomings.

In principle, it ought to be possible to adjust polls' confidence intervals based on their known limitations. And in fact, sophisticated pollsters and poll aggregators have long attempted precisely this (Silver, 2012). However, these adjustments are limited to known discrepancies between polls' samples and the voting population. Given that the list of discrepancies is long and changing, it seems inevitable that adjustment will be incomplete. How incomplete? Estimating that incompleteness requires specifying unknown unknowns, an epistemological impossibility. For instance, it is impossible to know what news stories might emerge between a poll and election day.

The approach we employ looks backward and measures the historical discrepancies between polls and elections. This can provide some guidance for estimating the size of such discrepancies. However, unless the past is perfectly representative of the future, it will provide an imperfect guide. And of course we live in a changing world whose future course is never perfectly predictable. The reasons why polls differ from elections in the future are not limited to the reasons they differed in the past.

We note three implications of our findings. First, authorities communicating the results of election polls should adjust downward the level of confidence they claim, or to adjust upward the size of the confidence intervals they report. If they wish to report well-calibrated confidence intervals, pollsters or journalists ought to substantially widen their confidence intervals. Current

practice, which reports confidence intervals assuming that the only source of error is sampling error, claims greater accuracy than polls provide. Second, consumers of polling data ought to have less confidence in poll results. Voters are making an error when they decide not to vote because they believe polling tells them how the election is going to come out. Nevertheless, we must note that data from our second study do not provide much reason for optimism with regard to voters' corrigibility. Even when we provided information about the poor historical performance of polls, respondents' confidence in their accuracy was insufficiently dampened. Third, we are skeptical of rules, such as those governing inclusion in political debates, which rely exclusively on polling data. Reliance on polls to decide which candidates deserve attention must be tempered with knowledge of the imperfections in polls' predictive accuracy.

Polls that fail to anticipate all the ways in which their results differ from the actual vote will be wrong for reasons that they do not expect. The result must be that they overestimate their accuracy and underestimate their propensity for error. Even pollsters who are willing to admit, in the abstract, that polls are vulnerable to various sources of inaccuracy beyond simple sampling error, may be reluctant to discount their own poll. After all, they have done all they can to minimize error and correct for known biases. The "inside view" of this current poll focuses on the benefits of correcting for known biases but it cannot account for unknown problems and novel issues (Kahneman & Lovallo, 1993). This dilemma may also partially explain why the lay participants in our second study resisted relying solely on historical base rates and instead expressed optimism that a new poll might be more accurate.

This is, of course, an epistemological conundrum shared by other statistical tools. Economists and sociologists have long been concerned about the degree to which "model uncertainty" impairs their inferences from samples (Hansen & Sargent, 2001; Young & Holsteen, 2017). Sometimes a scientific theory to explain some phenomenon is wrong and applies the wrong statistical model. But because scientists cannot test all possible models, they

cannot know whether their model of the data generating process is the right one and so they will come away overconfident in the value of their model.

Overconfidence is not unique to scientific theory testing or political polls. One of the most persistent biases in judgment is the excessive faith we know the truth, also known as overprecision (Moore et al., 2015). Because we base our beliefs on an imperfect and biased sample of information, sometimes we will be wrong for reasons that we did not anticipate. Our information sample may differ from the truth in systematic ways. And because it is difficult to consider the information we do not know, it is easy to overestimate the accuracy of our beliefs, including beliefs about elections. Our results suggest that this overprecision persists, not only in reports of poll results, but in people's interpretation of those reports.

References

- BBC News. (2016). *EU Referendum Results*. https://www.bbc.com/news/politics/eu_referendum/results
- Berlin, I. (1953). *The hedgehog and the fox: An essay on Tolstoy's view of history*. Simon & Schuster.
- Biemer, P. P. (2010). Total survey error: Design, implementation, and evaluation. *Public Opinion Quarterly*, 74(5), 817–848.
- Cohn, N. (2017, May 31). A 2016 Review: Why Key State Polls Were Wrong About Trump (Published 2017). *The New York Times*. <https://www.nytimes.com/2017/05/31/upshot/a-2016-review-why-key-state-polls-were-wrong-about-trump.html>
- Ford, R., Wlezien, C., Pickup, M., & Jennings, W. (2017). Opinion polling and vote. In K. Arzheimer, J. Evans, & M. S. Lewis-Beck (Eds.), *The SAGE Handbook of Electoral Behaviour* (pp. 787–812). SAGE Publications. <https://eprints.soton.ac.uk/412859/>
- Goot, M. (2021). How good are the polls? Australian election predictions, 1993–2019. *Australian Journal of Political Science*, 56(1), 35–55.
- Groves, R. M., & Lyberg, L. (2010). Total survey error: Past, present, and future. *Public Opinion Quarterly*, 74(5), 849–879.
- Hansen, L., & Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2), 60–66.
- Harvey, H. M. (2016, December 19). Skeptical 70,000 black voters abstained from presidential vote. *The Hill*. <https://thehill.com/blogs/pundits-blog/national-party-news/311099-skeptical-70000-black-voters-abstained-from>
- Isakov, M., & Kuriwaki, S. (2020). Towards Principled Unskewing: Viewing 2020 Election Polls Through a Corrective Lens from 2016. *Harvard Data Science Review*, 2(4). <https://doi.org/10.1162/99608f92.86a46f38>
- Jennings, W., & Wlezien, C. (2018). Election polling errors across time and space. *Nature Human Behaviour*, 2(4), 276–283. <https://doi.org/10.1038/s41562-018-0315-6>
- Jones Jr, R. J. (2008). The state of presidential election forecasting: The 2004 experience. *International Journal of Forecasting*, 24(2), 310–321.
- Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk

- and risk taking. *Management Science*, 39, 17–31.
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442.
- Lord, C. G., Lepper, M. R., & Preston, E. (1984). Considering the opposite: A corrective strategy for social judgment. *Journal of Personality and Social Psychology*, 47(6), 1231–1243.
- McMinn, S., & Hurt, A. (2020, April 21). *Tracking The Money Race Behind The Presidential Primary Campaign*. NPR.Org. <https://www.npr.org/2019/04/16/711812314/tracking-the-money-race-behind-the-presidential-campaign>
- Mercer, A., Deane, C., & McGeeney, K. (2016). *Why 2016 election polls missed their mark*. Pew Research Center. <https://www.pewresearch.org/fact-tank/2016/11/09/why-2016-election-polls-missed-their-mark/>
- Moore, D. A., Tenney, E. R., & Haran, U. (2015). Overprecision in judgment. In G. Wu & G. Keren (Eds.), *Handbook of Judgment and Decision Making* (pp. 182–212). Wiley.
- Moss, A., & Litman, L. (2020, June 12). *Demographics of People on Amazon Mechanical Turk*. CloudResearch. <https://www.cloudresearch.com/resources/blog/who-uses-amazon-mturk-2020-demographics/>
- Russonello, G. (2020, November 5). The Polls Underestimated Trump—Again. Nobody Agrees on Why. *The New York Times*. <https://www.nytimes.com/2020/11/04/us/politics/poll-results.html>
- Shirani-Mehr, H., Rothschild, D., Goel, S., & Gelman, A. (2018). Disentangling bias and variance in election polls. *Journal of the American Statistical Association*, 113(522), 607–614.
- Silver, N. (2012). *The signal and the noise: Why so many predictions fail—But some don't*. Penguin Press.
- Silver, N. (2016a, June 29). *2016 Election Forecast*. FiveThirtyEight. <https://projects.fivethirtyeight.com/2016-election-forecast/>
- Silver, N. (2016b, November 11). Why FiveThirtyEight Gave Trump A Better Chance Than Almost Anyone Else. *FiveThirtyEight*. <https://fivethirtyeight.com/features/why-fivethirtyeight-gave-trump-a-better-chance-than-almost-anyone-else/>
- Silver, N. (2020, November 3). *Biden's Favored In Our Final Presidential Forecast, But It's A Fine Line*

- Between A Landslide And A Nail-Biter*. <https://fivethirtyeight.com/features/final-2020-presidential-election-forecast/>
- Tetlock, P. E. (2005). *Expert political judgment: How good is it? How can we know?* Princeton University Press.
- Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Signal.
- The Commission on Presidential Debates. (2020). *CPD: Overview*. <https://www.debates.org/about-cpd/overview/>
- Voss, S., Gelman, A., & King, G. (1995). The polls—A review: Preelection survey methodology: Details from eight polling organizations, 1988 and 1992. *Public Opinion Quarterly*, 59(1), 98–132.
- Westwood, S. J., Messing, S., & Lelkes, Y. (2020). Projecting confidence: How the probabilistic horse race confuses and demobilizes the public. *The Journal of Politics*, 82(4), 1530–1544.
- YouGov. (2016, June 23). *YouGov on the day poll: Remain 52%, Leave 48%*. <https://yougov.co.uk/topics/politics/articles-reports/2016/06/23/yougov-day-poll>
- Young, C., & Holsteen, K. (2017). Model Uncertainty and Robustness: A Computational Framework for Multimodel Analysis. *Sociological Methods & Research*, 46(1), 3–40.