

Failure and success in political polling and election forecasting¹

Andrew Gelman

17 Aug 2021

1. A crisis in election polling

Polling got a black eye after the 2016 election, when Hillary Clinton was leading in the national polls and in key swing states but then narrowly lost in the Electoral College. The pre-election polls were again off in 2020, with Joe Biden steady at about 54% of the two-party vote during the campaign and comfortably ahead in all the swing states, but then only receiving a 52% vote share and winning some swing states by narrow margins.² The polls also overstated Democratic strength in congressional races. In other recent elections, the record of the polls has been mixed: they were accurate in the 2018 congressional elections,³ the Georgia Senate races in January,⁴ and recent British parliamentary votes,⁵ but were notoriously wrong on Brexit.⁶ A much-discussed aspect of recent pre-election polling errors has been an underrepresentation of choices favored by lower-education white voters—but pollsters are aware of this problem and, at least since 2016⁷ have been careful to adjust for ethnicity and education, yet this did not solve the problem of polls underestimating the Democrats' vote share in key swing states this past November.⁸ In particular, our own forecast with the Economist magazine (further discussed below) forecast Biden winning 54% of the two-party vote in Michigan and Wisconsin as compared to his actual vote share of 51% in each state, with similar errors in Florida, North Carolina, and other swing states.

There is more to political polling than election forecasting: pollsters inform us about opinion trends and policy preferences. But election surveys are the polls that get the most attention, and they have a special role in our discussions of polling because of the

moment of truth when poll-based forecasts are compared to election outcomes. The problem is challenging in part because pre-election polls attempt to survey a population—voters—that is undefined at the time of the sampling.

The recent successes and failures of pre-election polling invite several questions: Why did the polls get it wrong in some high-profile races? Conversely, how is it that the polls sometimes do so well? Should we be concerned about political biases of pollsters who themselves are part of the educated class? And what can we expect from polling in the future? The focus of the present article, however, is how it is that polls can perform so well, even given all the evident challenges of conducting and interpreting them. The key challenges are (a) attaining a representative sample of potential voters, and (b) predicting turnout. The technology now available to pollsters and the ability to develop ever more complex forecasting models with each election cycle will help, but we cannot expect to eliminate nonsampling errors, because conditions for new elections are always changing. Unlike sampling error, nonsampling error cannot be reduced simply by conducting more and larger surveys.

I come at these questions with the following background. As a statistician, I have worked on multilevel regression and poststratification (MRP), which is now a standard method used by YouGov and other pollsters for survey adjustment and small-area estimation (for example, to study opinion within states or congressional districts).⁹ As a political scientist, my colleagues and I have studied polling and public opinion, most notably finding that a large proportion of the swings in polls are due to differential nonresponse—that is, when a candidate is doing well, his or her supporters tend to be more likely to respond to surveys, thus causing a short-term feedback that magnifies small changes in opinion to misleadingly large swings in the polls.¹⁰ We saw strong evidence of this sort of polling bias in the 2012 and 2016 presidential elections, with

high correlations between each party's share of support in each poll and the percentage of their partisans among the respondents.¹¹ And as an unpaid collaborator with journalists at the Economist, I helped develop their 2020 presidential forecast, which at election eve gave Biden a 97% chance of winning, predicting that he would receive between 259 and 415 electoral votes and between 51.5% and 57.3% of the popular vote.¹² There are different ways this sort of forecast uncertainty can be conveyed.¹³

2. Survey errors and adjustments

Before discussing the successes and failures of political polls and forecasts, we briefly consider how they work. Opinion polling involves calling some number of people (or pulling them off an internet panel) and asking them several questions, including how they plan to vote in the upcoming election. There are several ways this can go wrong. The least important reason is sampling error: if you select 1,000 people at random, then it's unlikely your sample will exactly match the population in its political views. If you flip a coin 1,000 times, you might get 490 heads, or 515 heads, rather than exactly 500; similarly, if the population is 30% Democrats, 30% Republicans, and 40% Independents, you would not expect exactly 300, 300, and 400 in a random sample of 1000. A more important concern is nonresponse: the people who respond to surveys are not themselves a random sample of the general adult population or of the voters.¹⁴ This brings us to the challenge of turnout modeling: predicting who is likely to vote. Finally, a poll is a measurement at one point in time, not a forecast, and cannot tell us if people are responding insincerely or if they might change their opinion before election day.

Pollsters do their best to minimize nonsampling error in the analysis stage by adjusting for differences between sample and population. Various techniques are used,

depending on the goals of the adjustment. The simplest approach is weighting, where survey respondents are given numerical weights representing the number of people they are estimated to represent in the population; people in groups who are less likely to be included in the set of respondents (because of sample design or nonresponse) are given higher weights, and the weighted sample approximately represents the population. A separate, and also important, problem is estimating turnout by subgroup, because in an election forecast the goal is to adjust to the population of voters.

My own preferred approach of MRP works by dividing the problem into two parts. The first step is to train a statistical model (traditionally this is a multilevel regression, but more generally it could be a regularized prediction method from machine learning) on polling data to estimate voter turnout and the probability of supporting each candidate given demographics (age, sex, ethnicity, education), geography (in a national survey, this would be region, state, and perhaps urban/suburban/rural), and their interactions. The second step is to combine these estimates to get state and national estimates, weighting the estimates within each demographic/geographic segment by its population as obtained from the Census. For example, the estimate of Biden's support in Nevada is a weighted average of his estimated support in each demographic group in the state, weighting each group by its anticipated turnout in the election. There are several advantages of this approach. First, it produces inference for state-level opinions from data combined from multiple national polls ("partial pooling"); second, it gives estimates for small slices of the population ("small-area estimation"); third, it allows the modeling of both vote preference and turnout. The pollster Langer Research Associates demonstrated the use of MRP in the 2016 U.S. presidential election,¹⁵ and YouGov used the same approach to successfully predict constituency-level results in recent U.K. elections.¹⁶ Small-area estimation has value far beyond election

forecasting, as it can help us understand how people vote, for example in analyses of how the gender gap varies by age and education¹⁷ and distinguishing between mobilization (getting your supporters to vote) and persuasion (convincing undecided voters to support your candidate) in explaining electoral swings by breaking down estimated swings demographically.¹⁸

MRP has challenges too, both in statistical analysis and data collection. First, how exactly do you use a poll to perform small-area estimation? If we are estimating opinion in 4 age categories, 2 categories of sex, 4 ethnicity categories, 5 education levels, and 50 states, that's 8000 cells. Even if you are analyzing a set of polls with a total of 20,000 respondents, this will still leave with many cells with 0, 1, or 2 respondents, hardly enough to get any sort of estimate. Thus, any estimate at the cell level is necessarily model-based: indeed, multilevel regression is an approach to statistical modeling that performs this estimate by combining information from similar cells so that voting trends in Wyoming, for example, are estimated in part using data from other conservative western states. This approach performs well unless polling in these other states is off in a similar direction. The second difficulty of MRP is that we would like to adjust for non-census variables such as party identification and previous vote. This is an active area of research and involves modeling as well. The final challenge is that the adjustments only adjust for the variables included in the model. This perhaps is one reason that survey adjustments did better prediction for U.K. parliamentary elections (which are roughly characterized by uniform swing from the previous vote) than for Brexit (whose vote was cross-party and offered no easy baseline).

There is also the practical issue that any survey adjustment method (whether by weighting, MRP, or some other approach) requires data from the survey and the

population. But often all we have are the “toplines”: the summaries from the poll with no raw data and often with incomplete information on how the summaries were obtained. Indeed, for our Economist forecast we performed no MRP at all—no modeling and no poststratification. Instead, we just worked with the topline publicly released by the polling organizations, with the hope that individual pollsters did a good job of adjusting for nonresponse (and including in our model the possibility that they did not).

But, for all the difficulties of survey adjustment, we still do it, because we have no choice. Raw samples consistently look different from the general population of voters on key demographic and political variables such as sex, age, ethnicity, and party registration, so some adjustment needs to be done. Even for polls that use random digit dialing or some other method of probability sampling, nonrepresentativeness is a problem, which is no surprise given that response rates are typically under 10%.

3. Polls and election forecasting

Publicly forecasting a one-time event is risky for your reputation. On one hand, you want your forecast to be informative so you don’t want it unnecessarily vague; on the other hand, a precise forecast in the wrong direction is embarrassing. Our Economist forecast predicted national and state-level elections using a combination of historical election outcomes, economic and political “fundamentals” (economic statistics and presidential approval), and state and national polls.¹⁹ There were hundreds of available public polls during this campaign, and we accounted for systematic errors for each polling organization (“house effects”) and also shared errors affecting the mass of pollsters. Polling of humans is far from the simple random sampling described in many statistics textbooks, and indeed the experience of 2016 have made people generally

aware that national and state polls can be wrong at a level beyond what would be expected from sampling error. Our research based on state polls for president, senator, and governor in several past elections led us to conclude that nonsampling error is, on average about as large as sampling error in public polls, and this has impact on how to use polls in forecasting.²⁰ When constructing the Economist forecasting model, we allowed the polls' measurement of Biden's support to be off by roughly as much as 3 percentage points in either direction, which did cover the 2% error that occurred in the national popular vote as well as larger errors in some states—hence the prediction intervals that included the election outcome—but an error by that much was still on the high end of what we were anticipating.

But there is information other than polls. For months, online betting markets such as Betfair were anticipating a close election even while the polls showed a large Democratic lead: when poll-based forecasts such as ours were giving Biden an 80% or 90% chance of winning, online prediction markets had implied odds of 60% to 70%. Markets and polls aggregate different sorts of information, and recent research suggests that an average of the two can outperform either source alone.²¹

Some of the discrepancies between markets and polls can be attributed to hopeful bettors in a sluggish, low-dollar-value market where there was little incentive for arbitrage (taking advantage of incoherence in betting odds), but other aspects of the discrepancy can be explained by external factors specific to the 2020 election that were not included in the forecasting model; such factors included the chance of the coronavirus getting under control before the election vaccine, differences in turnout between the two parties with Republicans being more likely to go out and vote (this was anticipated but was difficult for pollsters to account for in their turnout models),

massive loss, invalidation, or controversy with mail ballots, or the Trump campaign using the courts or legislatures to cheat in some way.

4. Explanations for polling errors

At best, a poll is a measure of public opinion, telling us vote intentions. To some extent, polls can measure turnout—you can ask people if they have already voted, or if they plan to vote—but this is difficult because survey response rates themselves are so variable, and if a poll gets more respondents, these could very well include a greater or lesser proportion of nonvoters.

Why were the polls as imprecise as they were? Some possible explanations are differential nonresponse, differential turnout, changes in opinion, and insincere survey responses. These last two explanations, which can also be phrased as “last-minute swings” and “shy Trump voters,” are natural explanations, but I doubt they are a major part of the story in 2020. With political polarization as high as it’s ever been (for example as measured by the gap between Democrats and Republicans in presidential approval) there’s just not much opportunity for a last-minute swing, nor was there any notable last-minute news during the campaign as there was in 2016. As for Trump supporters who stated otherwise when polled: anecdotes aside, it seems doubtful there are many people who would go to the trouble of responding to a poll and then answering insincerely. For one thing, a Trump-supporting respondent who states a preference for Biden is actually hurting, not helping, his preferred candidate, in that Trump’s perceived unpopularity is one reason why some of his fellow Republicans would be less strongly supportive of him, and state-level polling errors in 2016 were not consistent with the “shy Trump voter” hypothesis,²² nor did analysis of 2020 polls support this theory.²³

Differential nonresponse and differential turnout seem like more plausible explanations of polling error. Surveys traditionally got too many respondents who were older, white, and well educated, and so they routinely adjust for discrepancies between sample and population in age, ethnicity, education, and other demographics.²⁴ Some polls also adjust for partisanship, using recorded party registration, stated party identification, or stated vote in the previous election. But even the surveys that made all these adjustments were off by about 2 percentage points, on average.²⁵ All that adjustment was not enough, which suggests that the survey respondents within these demographic and partisan categories were still not completely representative of the voting population. 50-70% of Americans who are eligible to vote in presidential elections do so,²⁶ but survey response rates have been declining for decades and remain under 10% for many political polls,²⁷ and some reports have found that socially isolated voters, who are less likely to participate in surveys, were more likely to be Trump supporters.²⁸ As for differential turnout, we discussed this earlier: 2020 was an unusual election, with an unprecedented high rate of early voting²⁹ followed by a high turnout of everyone else on election day, and the Republicans' edge in new-voter registration³⁰ could well have translated into a turnout lead not captured in the polls. An analysis of panel surveys by the American Association of Public Opinion Research attributes most of the polling error to differential nonresponse which varies geographically, with relatively more nonresponse by Republicans in strong Republican states.³¹

What about political biases? Are biased polls the reflection of a liberal bias among people who work in polling and the news media? There are reasons for skepticism about this claim. First, the claim of liberal bias is incoherent with the claim that Democratic support is overestimated because conservatives are lying to pollsters: if it

really is in the pollsters' political interest to bias surveys in the Democrats' direction, then why would conservative respondents be so eager to help them in that goal? Second, one of the most important assets of any pollster, whether commercial or nonprofit, is its reputation, and, for that, the clear motivation is to be as accurate as possible. The same sampling, interviewing, and adjustment techniques that can enable accurate election polling should work for business questions too; conversely, getting a major political race wrong is hardly an advertisement for your polling acumen. Third, polls are commissioned by different news organizations—various models, including that of the Economist and Fivethirtyeight.com, estimate “house effects” which find systematic differences between polling organizations. For example, Fox News surveys also overestimated Biden's support.³²

5. Where polls get it right

It's good to understand where and how polls got it wrong, but it's also important to ask why polls do so well.

Survey response rates in the U.S. and around the world are in decline,³³ there is widespread distrust of surveys in a way that is correlated with political attitudes, but... the polls are really not so far off. An error of 2 or 3 percentage points is a problem for predicting very close elections but otherwise is not so consequential. When the polls say 48% and the outcome is 52% (as with Brexit), that's a problem, even though the percentage error is not large. In the 2016 presidential election, the polls had Hillary Clinton with 53% of the two-party vote and she only received 51%, which was decisive because she would have needed about 52% of the national two-party vote to win the electoral vote.

The reason that polls are so accurate in percentage terms (even in the above well-publicized cases of misprediction) is that we know many of the key variables that predict vote—age, sex, education, ethnicity, previous vote, and state—and we know how to (imperfectly) adjust for these factors. In addition, there is little motivation for survey respondents to lie (as compared, for example, to surveys estimating socially unfavored or favored behaviors, where for example there has been a consistent overestimation of church attendance).³⁴

My point here is not to defend the polls—we do care about close elections, and these failures are real—but rather to note that their performance has been impressive. And for many purposes—indeed, just about anything other than a close election—it’s just fine to estimate opinion to within 4 percentage points, which is the upper bound of most national pre-election polling errors. For example, a recent Gallup poll found two-thirds of Americans supporting the legalization of marijuana.³⁵ Suppose this survey was in error, and support was actually as high as 70%, or as low as 60%. This could have some impact on lobbying and legislative decisions at the margin but would not change the general conclusion that legalization has clear majority support.

And here’s an example from sociologist David Weakliem of how opinion polling can help us understand an active political debate.³⁶ Weakliem quotes columnists in the New York Times and Wall Street Journal arguing that opposition to coronavirus restrictions come from “working-class people who are pushing back” and that calls to reopen businesses appealed to “breadwinners who can’t bus tables, process chickens, sell smoothies or clean hotel rooms over Zoom,” but they were “less compelling to college-educated suburbanites, who tend to trust experts and can work from home, watch their kids and spare a laptop for online kindergarten.” As Weakliem points out, this argument is not as clear as it might seem at first: “you can also think of reasons

that middle-class people might oppose restrictions. Middle-class jobs are more likely to allow some distance from co-workers and customers, for example, and middle-class people tend to go out more frequently for dining and entertainment. As a result, they might risk less and gain more from reopening. That's why we need data."

And, indeed, two Fox News surveys in October and an NPR survey in August showed no differences between attitudes of whites with and without college degrees, with about 40% of each group expressing the opinion that reopening the economy should be a priority. There are some differences by political party, but not so much when comparing different income and education levels, which does not fit into some punditry on the topic. Again, if these numbers were off by 4 percentage points, this would not appreciably change the story, as what is relevant is not the exact number or even whether a view is held by more than 50% of the population but rather the general finding that support for reopening is not strongly divided on class lines.

Another recent example of the value of polling came from direct surveys about responses to the coronavirus pandemic, which provided useful information about behavior that appeared to be accurate when compared to external data sources such as vaccination rates. These findings were used to tailor messages to encourage vaccination.

Finally, what of the future of polling and election forecasting? The experiences of the last two presidential elections may well make us wary of summarizing predictions using win probabilities. After the 2016 election, Nate Silver made the claim that a large part of the frustration expressed toward his Fivethirtyeight.com predictions was misplaced because readers did not really grasp the probability of the different outcomes. We may return to the traditional margin of uncertainty (supplemented to

account for nonsampling errors) or report inferences conditional on national errors as discussed above. In retrospect I wish that we'd expressed our 2020 Economist forecast conditionally: instead of simply stating a forecast interval and probability of each candidate winning, we could've graphed this interval as a function of the national polling error, thus indicating the confidence we had in our forecast at different levels of survey accuracy and making this dependence clear. Such a presentation could be easier to understand and also directly convey the relevance of measurement error in the forecast.

We can also expect to see discrepancies between public polling and the more intensive efforts by well-funded campaigns and advocacy groups, who can do more effective survey adjustment using the voter file, which has information including past turnout history on nearly 200 million Americans.³⁷ Aside from this, it would be good to see less focus on political campaigns and more on surveys of attitudes where there is no need for extreme precision and where we can learn about people's changing views on a range of policies³⁸—although this needs to be balanced by public demand for news in high-profile races, and tempered by the understanding that when surveying attitudes, there is no gold standard of comparison as there is with election outcomes. One piece of good news is that issue surveys, unlike election polling, does not require forecasting of voter turnout.³⁹ Politics is not just about who wins elections and reactions to their aftermath; it's also about policy, and public opinion is relevant to political negotiations and decision making.⁴⁰

¹ To appear in *Statistics and Public Policy*. We thank the editor and three reviewers for helpful comments.

² <https://projects.economist.com/us-2020-forecast/president>

³ Average polling error of less than 0.5 percentage points, according to Nate Cohn (1998), “What the Polls Got Right This Year, and Where They Went Wrong,” *New York Times*, <https://www.nytimes.com/2018/11/21/upshot/polls-2018-midterms-accuracy.html>

⁴ Average polling errors of 0.5 percentage points and 1.1 percentage points, according to https://www.realclearpolitics.com/epolls/2020/senate/ga/georgia_senate_runoff_election_perdue_vs_ossoff-7319.html and https://www.realclearpolitics.com/epolls/2020/senate/ga/georgia_senate_special_election_runoff_loeffler_vs_warnock-7318.html

⁵ Polling error of 1-2%, according to Anthony Wells (2019), “Initial analysis of YouGov’s election polling,” 13 Dec, <https://yougov.co.uk/topics/resources/articles-reports/2019/12/13/initial-analysis-yougovs-election-polling>

⁶ Polls showed 48%-46% in favor of Remain (<https://ig.ft.com/sites/brexit-polling/>) while the actual result was 52%-48% favoring Leave. This 4% error was large by historical standards and also missed the outcome.

⁷ Courtney Kennedy et al., “An evaluation of 2016 election polls in the U.S.: Ad hoc committee on 2016 election polling,” American Association for Public Opinion Research, <https://www.aapor.org/Education-Resources/Reports/An-Evaluation-of-2016-Election-Polls-in-the-U-S.aspx>

⁸ This essay focuses on U.S polls (with occasional references to the U.K.); in other countries, issues arise that we do not consider here, including from poor telephone and internet coverage, instability of party coalitions making elections less predictable, and electoral fraud and violence.

⁹ Ben Lauderdale (2019), “How YouGov’s 2019 General Election model works”, 27 Nov, <https://yougov.co.uk/topics/politics/articles-reports/2019/11/27/how-yougovs-2019-general-election-model-works>

Yair Ghitza and Andrew Gelman (2013), “Deep interactions with MRP: Election turnout and voting patterns among small electoral subgroups,” *American Journal of Political Science* 57, 762-776.

¹⁰ Andrew Gelman, Sharad Goel, Douglas Rivers, and David Rothschild (2016), “The mythical swing voter,” *Quarterly Journal of Political Science* 11, 103-130.

¹¹ Andrew Gelman and David Rothschild (2016), “Trump’s up 3! Clinton’s up 9! Why you shouldn’t be fooled by polling bounces,” *Slate*, <https://slate.com/news-and-politics/2016/08/dont-be-fooled-by-clinton-trump-polling-bounces.html>

¹² <https://projects.economist.com/us-2020-forecast/president>

¹³ Andrew Gelman, Jessica Hullman, Christopher Wlezien, and George Elliott Morris (2020), “Information, incentives, and goals in election forecasts,” *Judgment and Decision Making* 15, 863-880.

¹⁴ For an early review of the many differences between random samples and real-world populations, see D. Stephen Voss, Andrew Gelman, and Gary King (1995), Pre-election survey methodology: Details from nine polling organizations, 1988 and 1992. *Public Opinion Quarterly* 59, 98-132. There is no reason to think that problems of representativeness have diminished since then.

¹⁵ Chad P. de Jonge, Gary Langer, and Sofi Sinozich (2018), “Predicting state presidential election results using national tracking polls and multilevel regression with poststratification (MRP),” *Public Opinion Quarterly* 82, 419-446.

¹⁶ Chris Stokel-Walker (2019), “How YouGov became the UK's best but most controversial pollster,” <https://www.wired.co.uk/article/yougov-general-election-poll-mrp>

¹⁷ Robert Trangucci, Imad Ali, Andrew Gelman, and Douglas Rivers (2018), “Voting patterns in 2016: Exploration using multilevel regression and poststratification (MRP) on pre-election polls,” <https://arxiv.org/abs/1802.00842>

¹⁸ Yair Ghitza (2019), “Revisiting what happened in the 2018 election,” Catalist Analytics, https://medium.com/@yghitza_48326/revisiting-what-happened-in-the-2018-election-c532feb51c0

¹⁹ Drew A. Linzer (2013), “Dynamic Bayesian forecasting of presidential elections,” *Journal of the American Statistical Association* 108, 124-134. Merlin Heidemanns, Andrew Gelman, and Elliott Morris (2020), “An updated dynamic Bayesian forecasting model for the 2020 election,” *Harvard Data Science Review* 2(4).

²⁰ Houshmand Shirani-Mehr, David Rothschild, Sharad Goel, and Andrew Gelman (2016), “Disentangling bias and variance in election polls,” *Journal of the American Statistical Association* 113, 607-614. Michael Isakov and Shiro Kuriwaki (2020), “Towards principled unskewing: Viewing 2020 election polls through a corrective lens from 2016,” *Harvard Data Science Review*, 2(4).

²¹ Rajiv Sethi, Julie Seager, Emily Cai, Daniel M. Benjamin, and Fred Morstatter (2021), “Models, markets, and the forecasting of elections,” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3767544

²² Andrew Gelman and Julia Azari (2017), “19 things we learned from the 2016 election (with discussion),” *Statistics and Public Policy* 4.

²³ Josh Pasek, Kenneth Winneg, Kathleen Hall Jamieson, R. Lance Holbert, Daniel Romer, Bruce Hardy, Matthew Levendusky, Andrew Renninger, and Kate Kenski (2021), “Assessing polling errors in the 2020 U.S. presidential election: An examination of state-level panel data.”

<http://joshpasek.com/projects/annenber-iod/assessing-polling-errors-in-the-2020-u-s-presidential-election-an-examination-of-state-level-panel-data/>

²⁴ D. Stephen Voss, Andrew Gelman, and Gary King (1995), “Pre-election survey methodology: Details from nine polling organizations, 1988 and 1992,” *Public Opinion Quarterly* 59, 98-132.

²⁵ <https://projects.economist.com/us-2020-forecast/president>

²⁶ <http://www.electproject.org/home/voter-turnout/voter-turnout-data>

²⁷ Scott Keeter, Nick Hatley, Courtney Kennedy, and Arnold Lau (2017), “What low response rates mean for telephone surveys,” Pew Research Center, <https://assets.pewresearch.org/wp-content/uploads/sites/12/2017/05/12154630/RDD-Non-response-Full-Report.pdf>

²⁸ Daniel Cox (2020), “Could social alienation among some Trump supporters help explain why polls underestimated Trump again?”, <https://fivethirtyeight.com/features/could-social-alienation-among-some-trump-supporters-help-explain-why-polls-underestimated-trump-again/>

²⁹ <https://electproject.github.io/Early-Vote-2020G/index.html>

³⁰ Stephanie Saul (2020), “Behind in polls, Republicans see a silver lining in voter registrations,” *New York Times*, 19 Oct, <https://www.nytimes.com/2020/10/19/us/politics/vote-registration-Florida-North-Carolina-Pennsylvania.html>

³¹ Josh Pasek, Kenneth Winneg, Kathleen Hall Jamieson, R. Lance Holbert, Daniel Romer, Bruce Hardy, Matthew Levendusky, Andrew Renninger, and Kate Kenski (2021), “Assessing polling errors in the 2020 U.S. presidential election: An examination of state-level panel data.” <http://joshpasek.com/projects/annenber-iod/assessing-polling-errors-in-the-2020-u-s-presidential-election-an-examination-of-state-level-panel-data/>

³² For example, <https://www.foxnews.com/politics/poll-biden-trump-lead-narrows-slightly>

³³ Koen Beullens, Geert Loosveldt, Caroline Vandenplas, and Ineke Stoop (2018), “Response rates in the European Social Survey: Increasing, decreasing, or a matter of fieldwork efforts?” *Survey Methods: Insights from the Field*, <https://surveyinsights.org/?p=9673>

³⁴ C. Kirk Hadaway, Penny Long Marler, and Mark Chaves (1993), “What the polls don’t show: A closer look at U.S. church attendance,” *American Sociological Review* 58, 741-752.

³⁵ Megan Brenan (2020), “Support for legal marijuana inches up to new high of 68%,” *Gallup*, <https://news.gallup.com/poll/323582/support-legal-marijuana-inches-new-high.aspx>

³⁶ David Weakliem (2020), “Yes, ‘elites’ support coronavirus restrictions. So do working-class Americans,” *Washington Post*,

<https://www.washingtonpost.com/outlook/2020/12/01/coronavirus-restrictions-working-class-divide/>

³⁷ Yair Ghitza and Andrew Gelman (2020), “Voter registration databases and MRP: Toward the use of large scale databases in public opinion research,” *Political Analysis* 28, 507-531.

³⁸ Benjamin I. Page and Robert Y. Shapiro (1992), “The Rational Public: Fifty Years of Trends in Americans' Policy Preferences,” University of Chicago Press.

³⁹ Robert Y. Shapiro (2020), “Despite the 2020 election results, you can still trust polling. Mostly,” *Washington Post*, <https://www.washingtonpost.com/politics/2020/12/03/was-2020-really-disaster-polling/>

⁴⁰ Lawrence R. Jacobs and Robert Y. Shapiro (2001), “Politicians Don't Pander: Political Manipulation and the Loss of Democratic Responsiveness,” University of Chicago Press.