

The Real Story Of 2016

What reporters
— and lots of
data geeks, too
— missed
about the
election, and
what they're still
getting wrong.

By Nate Silver

On Friday at noon, a Category 5 political cyclone that few journalists saw coming will deposit Donald Trump atop the Capitol Building, where he'll be sworn in as the 45th president of the United States. It's tempting to use the inauguration as an excuse to finally close the chapter on the 2016 election and instead turn the page to the four years ahead. But for journalists, given the [exceptional challenges](#) that Trump poses to the press and the extraordinary moment he represents in American history, it's also imperative to learn from our experiences in covering Trump to date.

As editor-in-chief of FiveThirtyEight, which takes a different and more data-driven perspective than many news organizations, I don't claim to speak to every question about how to cover Trump. And I don't expect many of the answers to be obvious or easy. But in the part of the story that I know best, horse-race coverage,¹ the results of the learning process have been discouraging so far.

While data geeks and traditional journalists each made their share of mistakes when assessing Trump's chances during the campaign, their behavior since the

election has been different. After Trump's victory, the various academics and journalists who'd built models to estimate the election odds engaged in [detailed self-assessments](#) of [how their forecasts had performed](#). Not all of these assessments were *mea culpas* — [ours emphatically wasn't](#) (more about that in a moment) — but they at least grappled with the reality of what the models had said.²

By contrast, some traditional reporters and editors have built a revisionist history about how they covered Trump and why he won. Perhaps the biggest myth is when traditional journalists claim they weren't making predictions about the outcome. That may still largely be true for local reporters, but at the major national news outlets, campaign correspondents rarely stick to just-the-facts reporting (“Hillary Clinton held a rally in Des Moines today”). Instead, it's increasingly common for articles about the campaign to contain a mix of analysis and reporting and to make plenty of explicit and implicit predictions. (Usually, these take the form of authoritatively worded analytical claims about the race, such as declaring which states are in play in the Electoral College.) Furthermore, editors and reporters make judgments about the horse race in order to decide which stories to devote resources to and how to frame them for their readers: Go back and read their coverage and it's clear that The Washington Post was prepared for the possibility of a Trump victory in a way that The New York Times wasn't, for instance.

If almost everyone got the first draft of history wrong in 2016, perhaps there's still time to get the second draft right.

Another myth is that Trump's victory represented some sort of catastrophic failure for the polls. Trump outperformed his national polls by only 1 to 2 percentage points in losing the popular vote to Clinton, making them slightly *closer* to the mark than they were in 2012. Meanwhile, he beat his polls by only 2 to 3 percentage points in the average swing state.³ Certainly, there were individual pollsters that had some explaining to do, especially in Michigan, Wisconsin and Pennsylvania, where Trump beat his polls by a larger amount. But the result was not some sort of massive outlier; on the contrary, the polls were pretty much as accurate as they'd been, on average, since 1968.

Why, then, had so many people who covered the campaign been so confident of Clinton's chances? This is the question I've spent the past two to three months thinking about. It turns out to have some complicated answers, which is why it's taken some time to put this article together (and this is actually the introduction to a long series of articles on this question that we'll publish over the next few weeks). But the answers are potentially a lot more instructive for how to cover Trump's White House and future elections than the ones you'd get by simply blaming the polls for the failure to foresee the outcome. **They also suggest there are real shortcomings in how American politics are covered, including pervasive groupthink among media elites, an unhealthy obsession with the insider's view of politics, a lack of analytical rigor, a failure to appreciate uncertainty, a sluggishness to self-correct when new evidence contradicts pre-existing beliefs, and a narrow viewpoint that lacks perspective from the longer arc of American history.** Call me a curmudgeon, but I think we journalists ought to spend a few more moments thinking about these things before we endorse the cutely contrarian idea that Trump's presidency **might somehow be a good thing for the media.**

To be clear, if the polls themselves have gotten too much blame, then [misinterpretation and misreporting of the polls is a major part of the story.](#)

Throughout the campaign, the polls had hallmarks of high uncertainty, indicating a volatile election with large numbers of undecided voters. And at several key moments they'd also shown a *close* race. In the week leading up to Election Day, Clinton was only barely ahead in the states she'd need to secure 270 electoral votes. Traditional journalists, as I'll argue in this series of articles, mostly interpreted the polls as indicating extreme confidence in Clinton's chances, however.

So did many of the statistical models of the campaign, of course. While FiveThirtyEight's final "polls-only" forecast gave Trump a comparatively generous 3-in-10 chance (29 percent) of winning the Electoral College, it was somewhat outside the consensus, with some other forecasts showing Trump with less than a 1 in 100 shot. Those are radically different forecasts: one model put Trump's chances about 30 times higher than another, even though they were using basically the same data. Instead of serving as an indication of the challenges of poll interpretation, however, "the models" were often lumped together because they all showed Clinton favored, and they probably reinforced traditional reporters' confidence in Clinton's prospects.

But the overconfidence in Clinton's chances wasn't just because of the polls. National journalists usually interpreted conflicting and contradictory information as confirming their prior belief that Clinton would win. The most obvious error, given that Clinton [won the popular vote by more than 2.8 million votes](#), is that they frequently mistook Clinton's [weakness in the Electoral College](#) for being a strength. They also focused extensively on Clinton's potential gains with Hispanic voters, but less on indications of a decline in African-American turnout. At

moments when the polls showed the race tightening, meanwhile, reporters frequently focused on other factors, such as early voting and Democrats' supposedly superior turnout operation, as reasons that Clinton was all but assured of victory.

Post-election coverage has also sometimes misled readers about how stories were reported upon while the campaign was underway. The table below contains some important examples of this. Election [post-mortems](#) by major news organizations have tended to skirt past how much importance they attached to FBI Director James Comey's [letter to Congress](#) on Oct. 28, for instance, and [how much the polls shifted toward Trump](#) in the immediate aftermath of Comey's letter.

ELEMENT	PRE-ELECTION VIEW OF TRADITIONAL POLITICAL JOURNALISTS	POST-ELECTION VIEW OF TRADITIONAL POLITICAL JOURNALISTS	FIVETHIRTYEIGHT'S VIEW OF THE EMPIRICAL EVIDENCE (both before and after the election)
Interpretation of the polling	Polls proved that Clinton was almost certain to win.	Polls showed that Clinton was almost certain to win, and therefore they should never, ever be trusted again.	Polls showed an uncertain and volatile race with Clinton as a modest favorite and Trump with a real chance.
Role of Electoral College	The Electoral College provided a clear advantage to Clinton because of the "blue wall," the country's changing demographics, and her campaign's more sophisticated voter-targeting operations.	Obviously, the Electoral College was a problem for Clinton, but her underperformance was substantially the result of preventable tactical errors made by her campaign.	The Electoral College represented a major vulnerability for Clinton because her voters were disproportionately concentrated outside of swing states.
Proficiency of Clinton campaign	Clinton's savvy campaign had a better read on the election than public polls did, as evidenced by early voting data and internal polling. The campaign would give her a big Electoral College advantage in the event of a close popular vote.	Clinton's campaign was arrogant, misread the conditions on the ground, and had poor internal polling. It possibly blundered into costing their candidate the Electoral College by ignoring Michigan and Wisconsin until it was too late.	Campaigns probably don't have that much impact in presidential elections. Clinton's campaign made some tactical errors, but these likely weren't enough to cost her the Electoral College, especially given that she lost states such as Pennsylvania and Florida where she had campaigned extensively.
Impact of Comey letter	The Comey letter was a "game changer" of self-evident importance. But luckily for Clinton, it almost certainly wasn't going to be enough to change the outcome because her Electoral College position was so safe.	The Comey letter was a marginally important story that that the press covered proportionately. Comey can't be blamed for Clinton's loss because her lack of trustworthiness was well-established before the letter and the election shouldn't have been close in the first place.	The Comey letter corresponded with a sharp decline in Clinton's polling, of a large enough magnitude to change the election result. The media gave disproportionate attention to the Comey letter and other email-related stories.
Role of racial resentment	Racism, or "Racially charged" attitudes, may have helped Trump in the Republican primary, but they were likely a net disadvantage to him in the general election given the diversity of the electorate. There just weren't enough white voters to elect Trump.	Anger at the establishment and economic anxiety helped lead Trump to victory. "Racially charged" attitudes may have some correlation with these characteristics, but it demeans Trump voters and reveals liberal partisanship to attribute his victory to racism.	Empirically, racism is hard to untangle from other factors such as economic conditions and education levels. But the evidence is consistent with racism having had a significant impact on the campaign.

In other cases, the conventional wisdom has flip-flopped without journalists pausing to consider why they got the story wrong in the first place. For instance, it's now become fashionable to bash Clinton for having failed to devote enough resources to Michigan and Wisconsin. Never mind, for a moment, that these states wouldn't have been enough to change the overall result. (If Clinton had won Michigan and Wisconsin, she'd still have [only 258 electoral votes](#).⁴ To beat Trump, she'd have also needed a state such as Pennsylvania or Florida where she campaigned extensively.) The criticism is ironic given that many stories during the campaign [heralded the Clinton campaign's savviness](#), while skewering Trump for having campaigned in "[solidly blue](#)" states such as Michigan and Wisconsin. It's fair to question Clinton's approach, but it's also important to ask whether journalists put too much stock in the Clinton campaign's view of the race.

What exactly, then, is the "right" story for how Trump won the election? I obviously have a detailed perspective on this — but in a macroscopic view, the following elements seem essential:

- **First, the background conditions were pretty good for Trump.** Clinton was trying to win a third consecutive term for her party, replacing a fairly popular predecessor in President Obama, but she was doing that amid a mediocre economy and at a time of high partisanship. Various "fundamentals" models put together by political scientists and economists considered a matchup between a "generic" Republican and a "generic" Democrat (say, between Marco Rubio and Joe Biden) to be roughly a toss-up under these circumstances, or perhaps to slightly favor the GOP. While these models have significant limitations, they argue against the widespread presumption that the election was Clinton's to lose.
- **Second, demographics gave Trump a big advantage in the Electoral College.** Clinton won the popular vote by 2.1 percentage points, similar to George W. Bush's margin of victory over John Kerry in

2004, after which Bush claimed to have earned a mandate. But she lost in the biggest popular vote-versus-Electoral College discrepancy since 1876. Although Trump has protested otherwise, this discrepancy does not appear to have been mainly the result of tactical choices made by the campaigns. Instead it reflected demographics: White voters without college degrees, by far Trump's strongest demographic group, were disproportionately concentrated in swing states, while Clinton's coalition of minorities and college-educated whites (but with declining turnout among black voters) produced huge gains for her in states such as California and Texas without winning her any additional electoral votes.

- **Third, voter preferences varied substantially based on news events, and the news cycle ended on a downturn for Clinton.** As compared with recent presidential elections, there were a much higher number of undecided and third-party voters in 2016, probably because of the record-setting unpopularity of both Clinton and Trump. As a result, public opinion was sensitive to news coverage and events such as debates, with Clinton holding a national polling lead of as much as 6 to 8 percentage points over Trump in most of June, August and October, but Trump within striking distance in most of July, September and (crucially) November. Late-deciding voters broke strongly toward Trump in the final two weeks of the campaign, amid a news cycle dominated by discussion of the Comey letter and the WikiLeaks hack of Democratic emails.

This is an uncomfortable story for the mainstream American press. It mostly contradicts the way they covered the election while it was underway (when demographics were often assumed to provide Clinton with an Electoral College advantage, for instance). It puts a fair amount of emphasis on news events such as the Comey letter, which leads to questions about how those stories were covered. It's much easier to blame the polls for the failure to foresee the outcome, or the Clinton campaign for blowing a sure thing. But we think the evidence lines

up with our version of events. And if almost everyone got the first draft of history wrong in 2016, perhaps there's still time to get the second draft right.

I want to lay down a few ground rules for how this series of articles will proceed — but first, a few words about FiveThirtyEight's coverage of Trump. My view is that we had lots of problems, but that we got most of them out of the way good and early by [botching our assessment](#) of Trump's chances of winning the Republican primary. Among our mistakes: That forecast wasn't based on a statistical model, it [relied too heavily on a single theory](#) of the nomination campaign ("[The Party Decides](#)"), and it didn't adjust quickly enough when the evidence didn't fit our preconceptions about the race. Moreover, we "leaned into" this view in the tone and emphasis of our articles, which [often scolded the media for overrating Trump's chances](#). While it's challenging to judge a probabilistic forecast on the basis of a single outcome, we have no doubt that we got the Republican primary "wrong."

Something like the opposite was true in the general election, in our view. While our model almost never⁵ had Trump as an outright favorite, it gave him a much better chance than other statistical models, some of which had him with as little as a 1 percent chance of victory. [Independent evaluations](#) also judged FiveThirtyEight's forecast to be the most accurate (or perhaps better put, the least inaccurate) of the models. The tone and emphasis of our coverage drew attention to the uncertainty in the outcome and to factors such as Clinton's weak position in the Electoral College, since we felt these were misrepresented and neglected subjects. We even got into a couple of [very public screaming matches](#) with people who we thought were unjustly overconfident in Trump's chances.

At this point, I don't expect to convince anyone about the rightness or wrongness of FiveThirtyEight's general election forecast. To some of you, a forecast that showed Trump with about a 30 percent chance of winning when the consensus view was that his chances were around 15 percent⁶ will self-evidently seem smart. To others, it will seem foolish. But for better or worse, what we're saying here isn't just hindsight bias. If you [go back and check our coverage](#), you'll see that most of these points are things that FiveThirtyEight (and sometimes also other data-friendly news sites) raised throughout the campaign.

With that in mind, here's ground rule No. 1: These articles will focus on the general election. That's because we spent a lot of time last spring and summer reflecting on the nomination campaign. You can find our self-critique of our primary coverage [here](#). For other detailed reflections, I'd recommend my colleague Clare Malone's [piece](#) on what Trump's win in the primary told us about the Republican Party, and my [article](#) on how the media covered Trump during the nomination process.

Ground rule No. 2: These articles will *mostly* critique how conventional horse-race journalism assessed the election, although with several exceptions. The focus on conventional journalism in this article is not meant to imply that data journalists got everything right, however. There's obviously a lot to criticize in how certain statistical models were designed, for instance. But we've already covered these modeling issues at length both [before](#) and [after](#) the election, so I won't dwell on them quite as much here. As a quick review, however, the main reasons that some of the models underestimated Trump's chances are as follows:

- **Most of the models underestimated the extent to which polling errors were correlated from state to state. If Clinton were going to underperform her polls in Pennsylvania, for instance, she was also likely**

to do so in demographically similar states such as Wisconsin and Michigan.

- Several of the models were too slow to recognize meaningful shifts in the polls, such as the one that occurred after the Comey letter on Oct. 28.
- Most of the models didn't account for the additional uncertainty added by the large number of undecided and third-party voters, a factor that allowed Trump to catch up to and surpass Clinton in states such as Michigan.
- Some of the models were based only on the past few elections, ignoring earlier years, such as 1980, when the polling had been way off.

Put a pin in these points because they'll come up again. Interestingly enough, the analytical errors made by reporters covering the campaign often mirrored those made by the modelers. I'd also argue that data journalists are increasingly making some of the same *non-analytical* errors as traditional journalists, such as using social media in a way that tends to suppress reasonable dissenting opinion.

One final ground rule: The [corpus](#) for this critique will be The New York Times. Specifically, it will be stories published by the Times's political desk (as opposed to by its investigations team, in its editorial pages or by its data-oriented subsite, The Upshot). This is not an arbitrary choice. The Times, which hosted FiveThirtyEight from 2010 to 2013, is one of the two most influential outlets for American political news, along with The Washington Post. But also, the Times is a good place to look for where coverage went wrong. Few major news organizations conveyed more confidence in Clinton's chances or built more of their coverage around the presumption that she'd become the 45th president. (At one point, the Times actually referred to Clinton's "[administration-in-waiting](#)"). Articles commissioned by the Times's political desk regularly asserted that the Electoral College was a strength for Clinton, when in fact it was a weakness. Its

reporters were dismissive about the [impact of white voters without college degrees](#) — the group that swung the election to Trump. And the Times, like the Clinton campaign, largely ignored Michigan and Wisconsin.

Related: POLITICS PODCAST

Politics Podcast: Trump Vs. The Polls -- Vol. II

As you read these, keep in mind this is mostly intended as a critique of 2016 coverage in general, using The New York Times as an example, as opposed to a critique of the Times in particular. Most of these mistakes were replicated by other mainstream news organizations, and also often by empirically minded journalists and model-builders. At the same time, a relatively small group of journalists and news organizations, including the Times, has a disproportionate amount of influence on how political events are understood by large segments of the American public. (Media consolidation may itself be a part of the reason that Trump's chances were underestimated, insofar as it contributed to groupthink about his chances.) I think it's important to single out examples of better and worse coverage, as opposed to presuming that news organizations didn't have any choice in how they portrayed the race, or bashing "the media" at large. Obviously, I'm mostly taking a critical focus here, but in the footnotes you can find a list of examples of outstanding horse-race stories — articles that sagely used reporting and analysis to scrutinize the conventional wisdom that Clinton was the inevitable winner.⁷

So here's how we'll proceed. I've clipped a number of representative snippets from the Times's coverage of the campaign from the conventions onward. Each one will form the basis for a short article that reveals what I view as a significant error in how 2016 was covered. We're currently planning on about a dozen of

these articles — the idea is to be comprehensive — grouped into two broad categories. The first half will cover what I view as technical errors, while the second half will fall under the heading of journalistic errors and [cognitive biases](#). It's a somewhat fuzzy distinction, but important for what lessons might be drawn from them. The technical errors ought to be easier to fix, but they have narrower applications.⁸ The cognitive biases reflect more deep-seated problems and have more implications for how Trump's presidency will be covered; they're also the root cause of some of the technical errors. But they won't be easy to correct unless journalists' incentives or the culture of political journalism change. We'll release these a couple of articles at a time over the course of the next few weeks, adding links as we go along. Then I'll have some concluding thoughts. It's going to be a lot of 2016, at the same time we're also covering what's sure to be a tumultuous 2017. But the election is too important a story for journalists to just shrug and move on from — or worse, to perpetuate myths that don't reflect the reality of how history unfolded.