




Why (Almost) Everyone Was Wrong

The results of this year's presidential election made a mockery of analytical election forecast modelers.

By Pradeep Mutalik



Updated on November 9, 2016:

There is only one person who correctly forecast the U.S. presidential election of 2016. His name is not Nate Silver or Sam Wang or Nate Cohn. It is Donald Trump. Trump made a mockery of the predictions of all the erudite analytical election forecast modelers. Uttering the battle cry of “Brexit Plus,” he confidently grabbed the thin sliver of a chance that the models gave him by winning the Sun Belt states of Florida and North Carolina and then, in a near-miraculous example of threading the needle, flipping not just one but three of the ordinarily blue Rust Belt states that formed Hillary Clinton’s “firewall” — Wisconsin, Michigan and Pennsylvania — to red.

Like everyone else, I am stunned. In my pre-election *Abstractions* post below, I commented that the “science of election modeling still has a long way to go,” but I must admit that the distance is far

beyond what I had imagined. It seems pointless now to try to dissect the statewide predictions of the various models as I had promised to do — none of them were even remotely in the ballpark. It is unclear how long it will take before election forecasting is trusted again.

[abstractions]

You could be kind and say that the election results were not incompatible with the model that showed the most uncertainty (Nate Silver's), but there is no doubt that all the model builders completely missed the Trump win. His surprise victory took perfect advantage of the vagaries of the electoral vote system, even as the margin in the popular vote was razor-thin in favor of Clinton. But the modelers also missed something more fundamental, and they will have to revise their models to accommodate it. This was a [systemic undetected polling error](#) — a kind of invisible “dark matter” of polling — that underestimated support for Trump in key states by two to six percentage points.

In [a comment](#) on my original *Abstractions* post below, I said, “The uncertainty we have to account for is the uncertainty of things ‘we don’t know we don’t know.’” It turns out that this uncertainty is far larger than I thought. It will be interesting to read the full details of how Trump won as will no doubt be dissected in numerous post-mortems by both data analysts and pundits. In the meantime, it might be worth it to follow [Sean’s suggestion](#) and read about [the professor](#) who has been calling every single election correctly since 1982 on the basis of just 13 simple questions. He too, predicted a Trump victory, as did Michael Moore in [an article](#) that now seems prescient.

Originally posted on November 8, 2016:

Why Nate Silver and Sam Wang Are Wrong

Will the results of the U.S. presidential election discredit or vindicate popular election forecasting models?

As voters head to the polls, two of the most celebrated and successful election forecasters, Sam Wang of the [Princeton Election Consortium](#) (PEC), and Nate Silver of [FiveThirtyEight](#) disagree “bigly.” The former predicts the likelihood of a Clinton win at over 99 percent while the latter put the chance of a Clinton win in the mid to high 60s all of yesterday, rising to the low 70s today. This has resulted in Silver blasting a Huffington Post [article](#), which criticized his low estimate, and calling very high estimates “not defensible” in an invective-laced [twitter storm](#), and Wang responding with [a defense](#) of his 99 percent model. It was bound to happen: These two models represent the extremes in public prognostication. Most other models, including those of [The New York Times Upshot](#) and betting sites, rate Clinton’s chances somewhere in the middle, at about 85 percent.

Interestingly, both the PEC and FiveThirtyEight models agree that the aggregated polls show Clinton ahead of Trump by about 3 percent nationally, and their predictions for the number of electoral votes Clinton will get are 307 and 302, respectively (270 are needed to win). This disparity in the probabilities, and relative agreement in the number of electoral votes, validates my comments in the [last Insights column](#) — that aggregating poll results accurately and assigning a probability estimate to the win are completely different problems. Forecasters do the former pretty well, but the science of election modeling still has a long way to go.

The problem is twofold: First, modelers do not estimate a margin of error for their uncertainty, and second, there is far too little empirical data to validate probability models. Mathematically, there are standard ways to generate a probability number as Sam Wang [demonstrates](#). You let your model

run, generate the difference between polling percentages (about 3 percent), calculate the expected error (say, 0.8 percent) and out pops the probability (say, 92 percent). This is based on the premise that all the assumptions that went into the model are true (Sam Wang considers the median of recent polls to be the true number). However, to get the margin of error of this probability figure we need to assign weights to all the other assumptions that you did not make — after all, an assumption is a plausible but unproven proposition and therefore its opposite also has some probability of being true. You would need to calculate the probabilities of a Clinton victory in hundreds of alternate models in order to find the margin of error of the probabilities. While this meta-modeling would put the probability estimate in perspective and be more accurate, notice that in the absence of enough empirical data, the likelihood of alternate assumptions would still be arbitrary. True accuracy would require a complex model that incorporated many more features than current models do, using data from hundreds of presidential elections, and we don't have that luxury. An aggregation of existing models is the best simulation of such a meta-model that we have today.

There is historic evidence against Wang's tight forecasts — he predicted a gain of 53 plus or minus 2 house seats for the Republicans in 2010, and they actually gained 63, as Harry Enten of FiveThirtyEight [pointed out](#). On the other hand, there is no evidence that there is as much uncertainty as the FiveThirtyEight model suggests — it is far more likely that the observed tightening in the polls at this time is caused by people returning to their “home” party, as Wang has said. The best we can do is to aggregate models and discard the outliers — FiveThirtyEight and PEC — just as the modelers aggregate polls and discard outliers there.

Tomorrow, the results will either validate election forecasting models or show that this is a fledgling, imprecise science. Here at [Abstractions](#), we'll compare the detailed state predictions and try to determine which poll aggregation model was the best.

No matter which candidate wins, there is something to look forward to.