

# The Signal and the Noise

*Why So Many Predictions Fail—but Some Don't*

NATE SILVER



PENGUIN BOOKS

PENGUIN BOOKS

An imprint of Penguin Random House LLC  
penguinrandomhouse.com

First published in the United States of America by Penguin Press,  
an imprint of Penguin Random House LLC, 2012  
Published with a new preface in Penguin Books 2020

Copyright © 2012, 2020 by Nate Silver

Penguin supports copyright. Copyright fuels creativity, encourages diverse voices, promotes free speech, and creates a vibrant culture. Thank you for buying an authorized edition of this book and for complying with copyright laws by not reproducing, scanning, or distributing any part of it in any form without permission. You are supporting writers and allowing Penguin to continue to publish books for every reader.

Illustrations credits:

- Figure 4-2: Courtesy of Dr. Tim Parker, University of Oxford  
Figure 7-1: From "1918 Influenza: The Mother of All Pandemics" by Jeffery Taubenberger and David Morens, *Emerging Infectious Disease Journal*, vol. 12, no. 1, January 2006, Centers for Disease Control and Prevention  
Figures 9-2, 9-3C, 9-4, 9-5, 9-6, and 9-7: by Cburnett, Wikimedia Commons  
Figure 12-2: Courtesy of Dr. J. Scott Armstrong, The Warton School, University of Pennsylvania

THE LIBRARY OF CONGRESS HAS CATALOGED THE HARDCOVER EDITION AS FOLLOWS:

Silver, Nate

The signal and the noise : why so many predictions fail—but some don't / Nate Silver  
p. cm.

Includes bibliographic references and index.

ISBN 978-1-59420-411-1 (hc.)

ISBN 978-0-14-312508-2 (pbk.)

Special Markets ISBN: 978-1-52470-463-6

1. Forecasting. 2. Forecasting—Methodology. 3. Forecasting—history.  
4. Bayesian statistical decision theory. 5. Knowledge, Theory of. I. Title.

CBI58.S54 2012

519.5'42—dc23 2012027308

Printed in the United States of America

17 19 20 18 16

DESIGNED BY AMANDA DEWEY

While the author has made every effort to provide accurate telephone numbers, Internet addresses, and other contact information at the time of publication, neither the publisher nor the author assumes any responsibility for errors or for changes that occur after publication. Further, publisher does not have any control over and does not assume any responsibility for author or third-party Web sites or their content.

## PREFACE TO THE 2020 EDITION

**W**e were warned. And it didn't really make much difference. In 2015, Bill Gates said that a pandemic was more likely to kill millions of people than a war.<sup>1</sup> In 2016, after Donald Trump won the presidency—an event about which we'll have more to say later—President Obama's transition team left behind a “playbook” dealing with responses to a number of pandemic scenarios, including a novel coronavirus.<sup>2</sup> In 2019, Dr. Anthony Fauci cited a respiratory pandemic like SARS—that is, something exactly like COVID-19—in response to a reporter's question of what kept him up at night.<sup>3</sup>

To be clear, these predictions weren't terribly *specific*. Nobody said, “A novel coronavirus will originate in winter of 2019–20 in Wuhan, China.” None of these people were clairvoyant.

But they attested that the background risk was fairly high. On average, there has been a severe pandemic around every fifty years in the U.S.—the last two major ones before COVID-19 were influenza pandemics in 1918–19 and

1957–58—which might imply around a 2 percent annual chance. If such a pandemic killed the same proportion of Americans as died in the 1918–19 pandemic, or around 0.7 percent of the population, that would result in 2.3 million deaths.<sup>4</sup> Take 2 percent of 2.3 million and you wind up with 46,000 deaths per year, or more than ten times as many people as died in the September 11 attacks. That’s why people like Gates and Fauci were worried.

It’s harder to say exactly when COVID-19 *specifically* became such an acute threat. As late as mid-February, a “superforecasting” panel—a project led by Phil Tetlock, whom you’ll meet in chapter 2, and which generally has made successful forecasts—assigned only a 3 percent chance to the pandemic being as bad as it would turn out to be a month later.<sup>5</sup> And there were plenty of experts—even Dr. Fauci!—who downplayed the threat to the U.S. at times.<sup>6</sup>

But as outbreaks began to flare in Italy and other countries, the situation was getting worse quickly. By February 25, the Centers for Disease Control and Prevention had ominously warned Americans that they should prepare for “significant disruptions in their daily lives.”<sup>7</sup> And epidemiologists like Harvard’s Marc Lipsitch were predicting that 40 to 70 percent of the world’s population would eventually be infected with the coronavirus.<sup>8</sup>

And yet, it would be almost four weeks before California became the first state to issue a statewide stay-at-home order on March 19. I’m sorry for the cliché, but those days literally felt like the calm before the storm. You find yourself watching the storm clouds gather, hearing the first claps of thunder, recognizing the inevitability of what’s about to come. But you don’t know how bad it’s going to be—and you feel helpless to do anything about it.

I remember going to dinner at the MIT Sloan Sports Analytics Conference in Boston in early March with a smart group of people—sports executives, star athletes, academics, venture capital investors, a U.S. congressman, and several of the people you’ll meet in this book. It would have been ironic if these people—most of whom, in some form or another, are paid to assess risk and uncertainty—had been ignorant of the coronavirus threat. But we were pretty worried. We all had the sense that this was the last time in a while that there might be this sort of get-together. Did any of us actually do anything about it, though? Not really. (It’s not even clear that the conference should have been

held in the first place; a Biogen conference in Boston a week earlier had turned out to be a superspreader event.<sup>9</sup>)

I remember returning to New York City after the conference, and Mayor Bill de Blasio telling people to avoid using the subway.<sup>10</sup> Those of you who don't live in New York might not grasp quite how vital the subway is; telling a New Yorker to avoid the subway is like telling a fish to avoid swimming. And yet, offices, schools, bars, and restaurants were still open; there were still basketball and hockey games at Madison Square Garden; and crowds of tourists still thronged every major landmark in the city.

So on balance, the coronavirus was less a failure of *prediction*\* and more a failure to act upon expert guidance that, although it got a lot of details wrong and came with plenty of uncertainty, also got most of the big, directional stuff right.

Most of this book deals with why expert predictions often go wrong, along with what can make for more reliable forecasts. But since *The Signal and the Noise* was published in 2012, there have been an increasing number of cases like these, where experts actually posited a fairly high probability of world-changing events, only for these forecasts to be largely ignored or misunderstood.

In other words, making a good forecast is one thing and convincing people to take it seriously is another. Especially if the chance is under 50 percent but meaningfully higher than zero. Like, say, a 28.6 percent chance that Trump would be elected president.

I don't remember that much from Election Day. I remember voting, with voters in my super-blue district in Chelsea, Manhattan, tossing nods and smiles to one another, perhaps thinking they were electing the first woman president. I remember getting into the FiveThirtyEight office at 5 P.M.—it's a long night, so you want to take the morning off—and seeing exit polls that showed Hillary Clinton way ahead. While our model had Clinton favored (a 72.4 percent chance for her; a 28.6 percent chance for Trump) we had spent a lot of time

---

\* See also Slate Star Codex's Scott Alexander on this topic, from which this section draws inspiration <https://slatestarcodex.com/2020/04/14/a-failure-but-not-of-prediction/>. Note: The blog was (at least temporarily) deleted on June 22, 2020.

trying to convince people that Trump really had a shot, that the 28.6 percent was real. Maybe it had been much ado about nothing after all.

Then I remember, at some point around nine o'clock, the *New York Times's* Upshot needle starting to jitter toward Trump. In one swing state after another—Florida! Wisconsin! Michigan!

Clinton's numbers were coming up *just a little bit short* of where she needed to be. I remember how the office was very quiet, how time seemed to slow down between when it became obvious that Clinton was *probably* going to lose and when the election was formally called for Trump, at around 2:30 in the morning.<sup>11</sup>

Frankly, I remember feeling prescient when I realized that Clinton was likely to win the popular vote despite losing the Electoral College. It was precisely the sort of scenario we had been warning about. No, our model hadn't called the outcome exactly. But our model doesn't make "calls"; it deals in probabilities, and it had assigned a much higher probability to this sort of outcome—Trump winning the Electoral College because white working-class voters broke toward him in the Midwest—than other forecasts did.<sup>12</sup>

I remember driving up to Hyde Park, New York, on a cold, dreary Sunday after the election, trying to gain some perspective, and visiting the FDR National Historic Site. I remember wondering if the U.S. was in for a long, tumultuous decade, like it had been in the Great Depression and through World War II.

I spent weeks rereading mainstream media coverage of the election in the *New York Times* and other outlets for a series of essays on FiveThirtyEight that would later become "The Real Story of 2016."<sup>13</sup> Even though the polls were actually fairly tight toward the end of the race, any sign that Trump might win—like, say, that he was leading in polls of traditional swing states such as Ohio and Iowa—was dismissed by most of the press. If you'd read the *Times's* coverage, for instance, which literally referred to Clinton's "administration-in-waiting," you wouldn't have thought Trump's chances were 28.6 percent—you'd have thought they were zero.<sup>14</sup> And yet, these same outlets were quick to blame "the data" for their own lack of contemplation that Trump might somehow win when the data showed such an outcome was entirely possible.<sup>15</sup>

That's *not* to say everything had gone right during the election. The campaign had been a colossal shitshow from the moment Donald Trump descended the escalator at Trump Tower in June 2015 to announce his candidacy. Less

than a year later, Trump won the Republican nomination in an outcome that I initially—and publicly and sometimes obnoxiously—had been skeptical about.<sup>16</sup> Why? The basic theory was that because the Republican Party establishment had opposed Trump, he was likely to fail. The establishment usually won and got to decide on its preferred candidate—say, Mitt Romney in 2012.

But part of what this book tries to teach is to avoid quick-and-dirty heuristics and instead to apply a more rigorous and disciplined approach. While the Party Decides theory had some empirical backing, the sample size wasn't very large and there were a fair number of counterexamples.<sup>17,18</sup> Sure, the establishment *usually* wins—but what if an antiestablishment candidate like Trump had been ahead in the polls for months on end? And what if there were signs that the GOP establishment was in disarray?

It's not obvious how to weigh factors like these against one another, which is why it can be useful to build a probabilistic statistical model—or, failing that, to at least be fairly modest in your claims given the uncertainties involved. But I hadn't really followed my own advice. We never designed a statistical model to simulate the entire primary process; if we had (as we did in 2020), my view might have been different. Basically, I just kind of winged it and was sloppy, I found myself getting caught up in my initial, not-very-well-thought-out position and not wanting to back down. It was painful to watch Trump clear one hurdle at a time on his way to the nomination. And it was a reminder that good forecasts need to be built on a solid foundation or they're not going to be better than those of the average TV pundit.

The general election was a different story. We had our model up and running right away. As the polls shifted, at various points it gave Trump anywhere from a 12 percent chance to a 50 percent chance of winning, settling in at 28.6 percent in our final Election Day forecast. This is a decently high probability, about the chance of Boston Red Sox slugger David Ortiz (lifetime batting average of .286) getting a base hit.

Our model's odds were also generally much higher than the conventional wisdom about Trump's chances. Betting markets gave Trump somewhere between a 10 and 20 percent chance depending on which market you looked at, so if you'd followed our model's advice, you'd have bet on Trump and made a lot of money in the process. One of my ABC News colleagues, reflecting the

**Table 1: Trump Electoral College chances, final forecasts**

MODEL	TRUMP CHANCES
Princeton Election Consortium	<1%
HuffPost Pollster	2%
Daily Kos Elections	8% █
PredictWise	11% █
NYT Upshot	15% █
FiveThirtyEight*	29% █

\* Based on the FiveThirtyEight "polls-only" forecast (28.7%). Our polls-plus\* forecast gave Trump a 28.3% chance instead.

SOURCE: NYT UPSHOT.

conventional wisdom in New York and Washington newsrooms, put Trump's chances at 5 percent.<sup>19</sup> And certain other statistical models gave Trump as little as a 1 or 2 percent chance.<sup>20</sup>

I know not everyone will remember things this way. But after our experience in the primaries—our not-very-rigorous overconfidence that Trump wouldn't win—we wanted to be *sure* people knew that Hillary Clinton *did not* have the race in the bag. It didn't change how we built the model (it was basically the same model as in 2008 and 2012). But it did change how we talked about it. So, we wrote headlines like "Yes, Donald Trump Has a Path to Victory," "How Trump Could Win the White House While Losing the Popular Vote," and "Trump Is Just a Normal Polling Error Behind Clinton."<sup>21</sup> The week before the election, I got in a huge, somewhat profane Twitter argument with *HuffPo's* Ryan Grim, who had written that I was "unskewing" the polls in Trump's direction and implied I was creating panic when there was no need for Democrats to worry.<sup>22</sup> (*The Huffington Post's* model had given Trump just a 2 percent chance.)

I won't take you on a huge technical detour about why our model was more bullish on Trump than others.\* But the most important factor was it recognized that the outcome in different states was correlated rather than independent, especially if the states were similar demographically. That is to say, it wasn't a

\* See this story if you're interested in more on this subject: <https://fivethirtyeight.com/features/why-fivethirtyeight-gave-trump-a-better-chance-than-almost-anyone-else/>.



surprise that Trump won Michigan *and* Wisconsin *and* Pennsylvania *and* Ohio *and* Iowa, all states that Obama had won in 2012. All of these states had plenty of white working-class voters. All Trump had to do to win the Electoral College was slightly outperform his polls among that one group. That's part of why our model gave Trump a considerably better chance of winning the Electoral College than the popular vote.

Still, these pleas fell on somewhat deaf ears. Despite his upset win in the primary—and despite the surprise of Brexit a few weeks later, another event that the media was shocked by even though polls showed a close race—it was hard to get other journalists to take Trump's chances seriously, even after polls considerably tightened in the wake of FBI Director James Comey's letter to Congress about Clinton's e-mails.<sup>23,24</sup> As the *New York Times's* Amy Chozick later said on our podcast, the conventional wisdom *just wasn't ready* to believe that *Hillary Clinton* could lose to *Donald Trump*.<sup>25</sup> Clinton had seemed like the president-in-waiting since 2007, while Trump had been “a guy who had spent the past decade as a reality TV star.”

To be clear, my claim is not that anyone should have predicted months ahead of time that COVID-19 was sure to become a global pandemic—or that Trump's victory was guaranteed. Rather, my claim is that these were *reasonably high probabilities* of highly consequential outcomes, and that these risks were largely ignored. Why?

It's easy to say “the public is innumerate” or “people don't understand probability.” But I don't think it's quite that simple. At least in the case of the election, polls found that the public *wasn't* especially complacent about a Clinton win. In the *Los Angeles Times's* final poll, for instance—which actually asked voters to estimate the probability of each candidate winning—the average voter gave Trump a 43 percent chance, considerably higher than the models or the media did.<sup>26</sup>

And I don't think it's necessarily a question of mathematical aptitude per

---

\* Chozick's book, *Chasing Hillary*, is full of these sorts of honest reflections on how the media covered 2016 and comes recommended by yours truly.

se. I've played a lot of poker in my life, and although there are too few women players, the people you'll encounter are otherwise fairly diverse in terms of age, race, personality, background, and their approach to the game. There are *some* players who fit the math nerd stereotype, but they come from all walks of life. And yet, almost all of them who have played the game for any length of time have fairly keen probabilistic intuition. They know *roughly* what chances are that their opponent has a made hand versus a bluff; they recognize patterns; they remember when they've been in similar situations in the past. Not all of them are winning players, but by dint of experience—it's painful when you keep playing badly and lose money as a result—they're much better at poker than the average person you'd bring in off the street.

So what are the barriers to better probabilistic thinking? I think they basically fall into two buckets.

## “Thinking Slow” Is Vital—but It's Hard Work

All right, maybe you've never played poker. So let's take another, perhaps more familiar example: driving to work or school. You might not look at this as an exercise in probabilistic reasoning. But think about all the fairly sophisticated calculations that you're making: which route is fastest; how other drivers and pedestrians are going to behave; whether you can beat the light; what the risk of a traffic ticket is if you're running late to a big meeting.

It might seem second nature to you. But it's actually a lot of work; it's taken decades to train very powerful computers to drive even semicompetently.\*

Daniel Kahneman's 2011 book *Thinking, Fast and Slow*—an inspiration for this book—describes two types of cognition. They're called . . . well, “thinking fast,” or System 1, and “thinking slow,” or System 2. System 1 is quick and instinctual; it governs how you'd react (as an experienced driver) if a dog looks as though it's going to dart into the street and you have to decide whether to slam

---

\* Imagine if you had to learn it from scratch again. When I started biking around New York City last summer, I remember how mentally taxing it was for the first month or so. Until you've gained experience, you have to pay a lot of attention in Manhattan traffic or you risk hurting yourself or others.

on the brakes. System 2 is deliberate and conscious. It's what you'd use when deciding how you'd write an e-mail to your boss to ask for a raise.

As I've said, when you gain experience with a difficult task, it can sometimes transition from System 2 to System 1. Despite being a math guy who will sometimes spend hours with computer simulations analyzing poker hands I played earlier, I'm actually very instinctive and play quickly when I'm in the middle of the game. But it's *only because I've done the work* to gain experience and to study the game that my "gut instinct" is fairly good and I can more easily detect when something feels off.

The problem comes when we try to apply System 1, intuitive thinking, to complicated problems with which we have little to no experience. Think about the coronavirus, for example. Even though the projections were scary—even though, if you forced yourself to use System 2 thinking and actually write down a prediction of how your life would change, you might look at overflowing hospitals in Italy and lockdowns in China and recognize the magnitude of what was happening—none of it necessarily felt all that salient or visceral because most people in Western societies have never experienced a pandemic as bad as COVID-19 in their lifetimes. Plus, exponential growth is hard to grasp. If coronavirus cases are doubling roughly every three days, as they were in some places in the early stages of the pandemic, then an initial cluster of ten cases will become around *ten thousand cases* within a month; our brains aren't used to numbers like those. So to System 1, the first couple of days of lockdown might have felt more like a snow day. By contrast, East Asian countries that had experience with SARS in 2002 and 2003 generally had a much more effective response to COVID-19.<sup>27</sup>

Or think about the presidential election. Elections are more common than once-in-a-century pandemics. But events that happen only once every four years are still relatively rare, so the memories can linger a long time. As people looked at polls in 2016, for instance, they may have remembered 2008 and 2012, when Obama got a scare at various points in the campaign (Romney nearly tied him in polls following the first presidential debate in 2012, for instance) but in the end prevailed fairly comfortably. Nor was there an especially good recent example of a polling underdog winning a presidential election.

If you forced people to deliberately think about what Trump's chances were—to use System 2—they may have given Trump a decently high probability,

as in the *Los Angeles Times* poll. But if they were acting instinctually and emotionally—and reading and watching coverage from journalists who were *also* acting instinctually and emotionally—a 28.6 percent chance might not have felt real or salient.\*

So, there are times when you need to slow down and *not* trust your gut. But there are things that can make this harder. One can be oversimplified media narratives about what is happening as a news event unfolds. I have a lot of sympathy for journalists—I *am* a journalist, among other things—but there's a reason that journalism is often called "the first draft of history." It's because the *first draft is usually a complete mess*. It gets a lot of things wrong, and this is almost impossible to avoid.

Developing news stories require a lot of probabilistic, provisional thinking. So when journalists or other experts are presenting them in too neat a package, be worried. Conversely, you should trust journalists and other experts *more* when they use phrases such as "I don't know," "I'm not sure," "we're still learning more," or "I've changed my mind."

The changing-your-mind part is especially important. Perhaps the hardest part of forecasting is knowing how quickly to revise your priors as new information is revealed. We'll have a lot more to say about priors in the second half of this book. You can easily err in either direction, saying "this is fine" as your original prediction is going up in flames—or flopping around like a weather vane. In the current media ecosystem, though, not changing your mind quickly enough is probably the bigger risk. As you become affiliated with a given position, you tend to form alliances around it, seeking positive reinforcement from others who think the same way—or readers may seek you out because they're looking to have their own priors confirmed. That can make switching more costly than it would be if you were solely concerned with making the most accurate forecast.

Also, resist the tendency to fight the last war. It's true that having even *one*

---

\* Frankly, I don't think looking at forecasts like FiveThirtyEight's changes most people's priors much either way, since elections are an emotional crescendo, and people are deep in System 1 mode by Election Day. If you're a Democratic voter, and System 1 is telling you to be worried, then a 28.6 percent chance will sound pretty high. (Would you board a flight with a 28.6 percent chance of crashing?) Conversely, if it's telling you to relax, then Clinton's 71.4 percent chance might sound comforting.

relevant example to work from can be a lot better than *zero*. (Think again about how countries with SARS experience did better against COVID-19.) But it can also be easy to overcompensate. There is a now-infamous clip (well, infamous to election geeks like me) of a group of pundits on MSNBC's *Morning Joe* on the morning before the Virginia gubernatorial race in 2017 in what was probably the highest-profile election since Trump's upset win in 2016.<sup>28</sup> Although Democrat Ralph Northam led by around three points in polls, which would translate into his winning around 75 percent of the time, the panelists were emphatically unwilling to predict a Northam victory, with some suggesting that Republican Ed Gillespie would prevail instead.<sup>29</sup> Their thinking had somehow morphed from (pre-Trump) "the candidate who is three points ahead in the polls *always* wins" to (post-Trump) "the candidate who is three points ahead in the polls *never* wins" without pausing to stop at "the candidate who is three points ahead in the polls *usually* wins," which was the right answer all along. As it turned out, Northam won by around nine points, a considerably bigger margin than polls predicted.

Most people fail to recognize how much easier it is to understand an event after the fact when you have all the evidence at your disposal. That's not to say every bad prediction should be excused. Far from it. The whole purpose of this book is to describe how society can make better predictions. But *making better first guesses* under conditions of uncertainty is an entirely different enterprise than second-guessing. Asking what a decision maker believed *given the information available to her at the time* is a better paradigm than pretending she should have been oracular.

## Beware the Herd

If I've changed my mind on one big theme of *The Signal and the Noise* since it was published in 2012, it's that I'm increasingly skeptical about the wisdom of crowds.

The book equivocates on this subject quite a bit. As I'll describe in chapter 11 about the efficiency of financial markets, for example, it's not that markets are perfectly efficient so much as it's awfully hard to beat them. Every time

you're making a trade—whether you're trading a stock or, say, a baseball player—there's another, probably equally smart and well-trained person on the other side of that trade. (Or to expand the metaphor further, suppose you're engaging in an *exchange of ideas*.) Are you sure you know something they don't? Are you sure *they* don't know something that *you* don't? If your side of the bargain is so appealing, then why are they willing to trade with you in the first place?

There is obviously some sort of balance here. On the one hand, it's presumptuous to think that you, just as one person, are wiser than the consensus around a subject, which—in theory, reflects a lot of smart people's views and experiences combined. On the other hand, society would never move forward if nobody questioned that consensus, and it's not as though the consensus has delivered us to a particularly good place. (Or at least it certainly doesn't seem that way as I write this amid a pandemic that has already killed more than 120,000 Americans as of June 2020.)

All I'm saying is that I think that balance has shifted, at least slightly. I think it's shifted in a way that may require you to trust the consensus a bit less and question the conventional wisdom a bit more. That's not to say you should be some sort of professional contrarian. Most people still overweight their own judgment. But I think you should be more skeptical than I would have said eight years ago.

James Surowiecki's 2004 book, *The Wisdom of Crowds*, is all about when collective judgment tends to be better—or worse—than individual judgment. The “collective” can be in the form of an organization or business or an industry or a market. The point is that there is some mechanism to form a *consensus* that does not merely reflect the views of any one person in the group.

I'm simplifying and modifying Surowiecki's thesis slightly, but there are basically three characteristics that he says are associated with wise collective judgment. As you read these, think about how well or poorly these categories are going in the United States today.

1. *Diversity*. If everybody comes in thinking the same way about a problem, then forming a group is pointless. A room of one hundred Nates (a scary thought!) is going to come to pretty much exactly the same conclusion that one Nate does, only they'll probably have become more hardheaded and inflexible about

- it by constantly agreeing with one another. So you want people with a diversity of backgrounds, experiences, and skill sets as part of your group.\*
2. *Independence.* People need to be able to share ideas and dissent openly, without fear of reprisal. Otherwise, the most powerful members of a group will just cow everyone else into submission and create a false consensus. Moreover, you want to avoid information cascades wherein people's preferences are dependent on everyone else's preferences.
  3. *Trust.* Groups are effective when people are confident that they are fair and free of corruption. And in any sort of group that purports to be representative, such as a government, people need to trust that their representatives have the collective interest in mind rather than looking out for themselves.

Among these categories, I would claim American society is arguably getting better on the first one, but much worse on the other two dimensions.

The U.S. has made some strides on racial and gender **diversity**, although—as I write this amid the largest sustained antiracism protests since the 1960s—it obviously has a long way to go. The number of women in Congress has roughly doubled in the past twenty years, for instance (although it is still just 24 percent). Surowiecki is generally talking about diversity of thought or diversity of skill sets more than racial or gender diversity per se, but these things are related; rooms full of white men are going to have a lot of blind spots.

But institutions like the media may have become less representative in certain other respects. A 2018 study found that around half of employees at the *New York Times* and *Wall Street Journal* attended 1 of 29 elite universities, or 12 law schools or business schools, representing an almost literal 1 percent of the roughly 3,700 degree-granting institutions in the United States.<sup>30,31</sup> With education increasingly becoming (along with race) the key political dividing line in the country—white voters who didn't attend college were pivotal to Trump's victory in 2016—there's room for concern if people in the media and other institutions all attended elite colleges and all share a certain set of political views.<sup>32</sup>

And while I don't want to be one of those people who constantly complain

---

\* Surowiecki also talks about another dimension, decentralization—bringing local knowledge or experience to bear on a problem—although I really think of this as a subset of diversity.

about social media . . . I'm going to complain about social media, at least for a couple of paragraphs. Twitter and Facebook generally amplify consensus and drown out or pile on dissenting views.<sup>33</sup>

It makes **independence** a lot harder when you face a social sanction for delivering an unpopular opinion. (It also makes changing your mind harder, as I mentioned earlier.) Or you may never hear those views in the first place if you've filtered the conversation such that you're hearing only from people you agree with.

Although it didn't inspire as much self-flagellation in the media as Trump's primary and general election victories in 2016, Joe Biden's win in the 2020 Democratic primary reflected this. Biden's chances were constantly discounted even though (like Trump in the 2016 GOP primary) he led in polls almost wire to wire, at least until his loss in the Iowa caucus.<sup>34</sup> And yet, by mid-October, betting markets—usually a good proxy for the conventional wisdom about the race—gave Biden merely an 18 percent chance to win the nomination and Elizabeth Warren better than a 50 percent chance, even though Biden was (slightly) ahead of her in polls at the time.<sup>35</sup>

Why was this? It may have been because Warren voters were overrepresented on social media (and that journalists pay a *lot* of attention to social media). Democrats active on social media were far more likely to be liberal, college educated, and white, all characteristics that described Warren voters, whereas Biden's voters were older, more moderate, more working-class, and more likely to be black or Hispanic.<sup>36</sup>

Finally, on the dimension of **trust**, American institutions are in grave trouble. From Congress, to the church, to the media, to schools, to the White House, most major American institutions have suffered a precipitous decline in the number of people who say they have confidence in them, according to Gallup polling.\*

Which of these institutions *deserve* the loss of public confidence and which are scapegoats is another question. But either way, it creates a potential problem. It's not just that people may not trust what these groups say, but also that

---

\* And one of the few institutions whose numbers had held up relatively well as of Gallup's last round of polling in 2019, the police, saw their numbers decline sharply in the wake of the killing of George Floyd by a Minneapolis police officer and the subsequent protests.



TABLE 2: PERCENTAGE OF AMERICANS SAYING THEY HAVE A "GREAT DEAL" OR "QUITE A LOT" OF TRUST IN EACH INSTITUTION<sup>37</sup>

Institution	Peak	In 2019
Public schools	62% (1975)	29%
The presidency	71% (1991)	38%
The church/organized religion	68% (1975)	36%
Congress	42% (1973)	11%
Banks	60% (1979)	30%
Newspapers	51% (1979)	23%
Television news	46% (1993)	18%
The Supreme Court	56% (1985, 1988)	38%
The military	85% (1991)	73%
Big business	34% (1975)	23%
The police	64% (2004)	53%
Organized labor	39% (1977)	29%
The criminal justice system	34% (2004)	24%
The medical system	44% (2003, 2004)	36%
Small business	70% (2017)	68%

the feeling can be mutual and self-reinforcing, and these institutions may not trust the public.

Public health experts in the U.S. initially advised against the importance of wearing masks at the start of the COVID-19 pandemic, for example, a move that was inspired at least in part by concerns that asking people to wear masks could trigger a run on personal protective equipment needed by health-care workers.<sup>38</sup> The experts—maybe correctly and maybe not—didn't trust the public enough to say, "Masks are probably helpful, but make your own or wait a few weeks instead of bulk ordering them on Amazon."

And advice on social distancing was sometimes inconsistent, with mixed messaging on what was safe (Outdoor protests? Trump political rallies?) or what

was a worthwhile trade-off. As I write this in June, trust in medical experts has fallen significantly over the course of the pandemic, especially among Republicans.<sup>39</sup> There's a lot to untangle in the politics of the coronavirus, and most of it is probably best left for when the pandemic has eventually passed us. But it has reinforced the breakdown of trust between experts and the public.

So there are reasons to think, per Surowiecki's criteria, that group and collective judgment is becoming worse among major institutions in the United States: that we are experiencing fewer of the benefits of consensus and more of the downsides of groupthink. Unfortunately, that doesn't mean our individual judgment is getting any better. We should be concerned about our own biases and the other limitations of our judgment, *and* we should be wary of groupthink.

There's a lot to be said against the staid, monolithic, "both sides" journalism that was the dominant paradigm until a few years ago—and which proved not up to the task of contemplating Trump's candidacy or his presidency. I think the eclectic mix of journalistic modes and perspectives that have replaced it are superior, on balance. But it's now much more likely that the news you encounter will be customized to your particular preferences; it's like going from a society built on arranged marriages to one where everyone meets on Tinder. The signals that you're getting aren't the ones that you're used to and may flatter your sensibilities more than you realize.

So we should be particularly wary when we hear simplified stories about complex events that match our political preferences or validate our beliefs. A liberal doesn't need to be told to be skeptical of Fox News. But when it seems like all of your friends are saying the same thing—*that's* when you need to be on high alert. It doesn't mean those views are necessarily wrong. Hopefully, you'll have selected your news sources well (and your friends well).<sup>\*</sup> But if a news story stimulates an instinctual, System 1 reaction, as can be especially likely when it's shared on social media, slow down and make sure it checks out with System 2. Update your priors.

I've sometimes thought that if *The Signal and the Noise* could be condensed to a bumper sticker, it would read THINK PROBABILISTICALLY. To that I'd add now add a second one: SLOW DOWN AND BE SKEPTICAL.

---

<sup>\*</sup> I'm certainly not claiming that the *New York Times* and One America News Network are equivalent, for example.