

Spotting Bullshit

IN SEPTEMBER 2017, A SHOCKING PHOTOGRAPH MADE THE ROUNDS on social media. Set in the locker room of the Seattle Seahawks football team, the picture appeared to show former Seahawks defensive lineman Michael Bennett bare-chested and brandishing a burning American flag. Circled around Bennett and cheering jubilantly were his Seahawks teammates and coaching staff.

This photograph was, of course, a fake. No Seahawk ever burned an American flag during a team meeting—or any other time. The photo (minus the burning flag) had been taken nearly two years earlier, when the Seahawks were celebrating a crucial victory against the rival Arizona Cardinals. But it spoke volumes about an ongoing cultural battle that had swept up the National Football League. San Francisco quarterback Colin Kaepernick led a growing number of players across the league in taking a knee during the national anthem as a protest against police brutality. Donald Trump branded these players as unpatriotic, antimilitary, and altogether un-American. The image of Michael Bennett with the burning flag, first posted to the Facebook page of an organization called “Vets for Trump,” provided an extreme expression of Trump’s narrative. Many viewers overlooked the low-quality image manipulations and shared the photograph, with angry invectives directed at the Seahawks players. The anger and disgust that many felt in seeing the image likely overshadowed any inclination to critically evaluate its authenticity and primed them to fall for the bullshit.

If bullshit is everywhere, how can we avoid being taken in? We

think it is crucial to cultivate appropriate habits of mind. After all, our habits of mind keep us safe on a daily basis. We don't think about it, necessarily, but as we drive to work, our eyes are scanning for a driver about to run a red light. Walking alone at night, we are aware of our surroundings and alert for signs of danger. Spotting bullshit is the same. It takes continual practice, but with that practice one becomes adept at spotting misleading arguments and analysis. While developing a rigorous bullshit detector is a lifelong project, one can go a long way with a few simple tricks that we will introduce in this chapter.

1. QUESTION THE SOURCE OF INFORMATION

Journalists are trained to ask the following simple questions about any piece of information they encounter:

Who is telling me this?

How does he or she know it?

What is this person trying to sell me?

These questions are second nature to us under some circumstances. When you walk into a used-car dealership and the salesman starts talking about how the car in the corner of the lot had only a single owner, a little old lady who drove it once a week to church on Sunday, you are, of course, thinking this way: Who is this person? A used-car salesman! How does he know this? Well, maybe he heard it straight from the little old lady herself. Or maybe, he heard it from the dealer across town who sold him the car. Or, just maybe, there never was a little old lady to begin with. What's he trying to sell you? That one's obvious. The 2002 Pontiac Aztek you made the mistake of glancing toward as you walked on the lot.

When we scan through our social media feeds, or listen to the evening news, or read the latest magazine page about how to improve our health, we need to ask the same questions.

In the process of writing this chapter, we read online that crystals “retain all the information they have ever been exposed to. Crystals absorb information—whether a severe weather pattern, or the experi-

ence of an ancient ceremony—and pass it to anyone that comes into contact with them.” Now this doesn’t even remotely jibe with our understanding of physics, so it’s worth asking ourselves these three questions about this claim.

Answering the first question is relatively straightforward. Who is telling us this? This text comes from an interview about healing crystals from the website of the lifestyle brand Goop. The interviewee is Colleen McCann, a “fashion stylist turned energy practitioner” and “certified shamanic energy medicine practitioner” who “utilizes a combination of crystals, color theory, chakra systems, astrology, naturopathy, and Feng Shui principles.”

Answering the second question, “How does she know it?,” can be harder to ascertain. In this case, however, McCann’s website gives us enough material to make some educated guesses. In her bio, we learn that McCann “started hearing voices, also known as *Clairaudience*, in a Brooklyn bodega” and that the “reputable Manhattan psychic” she subsequently consulted “gently broke the mystical bomb to her that she too was psychic. The voices were in fact her Spirit Guides saying hello.” She then “jumped heels first into the crystal-laden rabbit hole with three years of private mentorship with an Intuitive Healer. Her journey then led her to train at The Four Winds Society in Peruvian Shamanic Studies.” Finally, we learn that she “has spent a decade studying with a Buddhist Feng Shui master to learn Crystal Healing and Space Clearing.” Perhaps it was through these experiences, or others like them, that McCann picked up the information that she shared in this interview.

And the third question, “What are they trying to sell us?” Here again we have to guess a little bit, but only a little bit. The Goop company and the interviewee may be selling slightly different things, of course. McCann may be aiming to sell us on a set of ideas or a philosophy. In addition, and perhaps not coincidentally, her website reveals that she also sells crystals and provides services including “intuitive crystal readings,” “crystal gridding,” and “crystal luxtraction.” The Goop company, for their part, might argue that they are promoting a lifestyle. But they also sell a so-called Goop Medicine Bag, wherein for eighty-five dollars one receives a set of eight “chakra-healing crystals.” To us these seem nearly indistinguishable from the polished gems and such that you would get for five dollars a bag at

tourist shops—but think again. The Goop stones have been “cleansed with sage, tuned with sound waves, activated with mantras, and blessed with Reiki.”

In short, people may be trying to sell you used cars or life insurance or beauty treatments—or they be trying to sell you ideas, viewpoints, and perspectives. Some sales jobs get you to part with your hard-earned money. Other sales jobs convince you to believe something that you didn’t believe before, or to do something you wouldn’t have done otherwise. Everyone is trying to sell you something; it is just a matter of figuring out what.

We could also ask these questions about the photograph of Michael Bennett with a burning flag. Who is telling me this? The Facebook group called Vets for Trump. How do they know it? Since the photo appeared only on Facebook and was not reported in any traditional media outlet, the only possible story would be that someone had a camera in a locker room, but somehow the media was either not present or they all agreed not to report on what had happened—and none of the Seahawks players or staff spoke up about it afterward. That seems highly implausible. What are they trying to sell us? They want to convince us that the NFL players who are protesting racial injustice hold anti-American sentiments and may be a threat to the country. Even without the giveaway signs of a poor photoshopping job, the answers to these three questions should be enough to make us question the authenticity of such a shocking and unexpected picture.

2. BEWARE OF UNFAIR COMPARISONS

“**A**irport Security Trays Carry More Germs Than Toilets!” Media outlets around the world ran some version of this headline after a research study was published in September 2018, confirming the fears of every germophobe who has ever suffered through the airport security screening process.

But the claim is somewhat disingenuous. The scientists who did this study looked only at respiratory viruses, the kind transmitted through the air or through droplets on people’s hands when they cough or sneeze. It isn’t surprising that security trays have more respiratory viruses than toilet seats. People don’t usually cough or sneeze onto toilet seats, nor do they tend to touch them extensively with

their hands. Toilet seats have plenty of germs, just not of the kinds that the researchers were tallying.

Airline trays may be important vectors for colds and flus, but when the headlines bring toilet seats into the picture, they are making an unfair comparison for shock value. Trays don't have more germs than toilet seats, they just have more germs of the type likely to land on trays.

Let's look at another example. People have always enjoyed ranked lists. In a clickstream economy, where advertising revenues depend on page views, they're gold. A single top-ten list can generate ten page views per reader by putting each list item on a separate page. Farewell Casey Kasem, hello "12 Reasons Why Sam, the Cat with Eyebrows, Should Be Your New Favorite Cat."

One form of list that continuously reappears is some variant of "America's Most Dangerous Cities." Recently we came across such a list, released by financial news outlet *24/7 Wall St.* and based on a compilation by the FBI. At the top of the list were

1. St. Louis, MO
2. Detroit, MI
3. Birmingham, AL
4. Memphis, TN
5. Milwaukee, WI

Well, that got personal quickly. Carl was born in St. Louis; he spent much of his later teens exploring the neighborhoods of Detroit. Are these cities really all that bad? he wondered. The rankings are based on hard numbers from a trusted government agency. But is this really a fair comparison? Or is there something funny going on that sends St. Louis and Detroit to the top of the list? The first question we could ask is how they have quantified how dangerous a city is. Biggest potholes? Most bedbugs? Greatest number of golfers hit by lightning?

In this case, the metric of danger is the number of violent crimes per capita. We might try to argue that this measure does not do a good job of characterizing how dangerous a city is. Perhaps violent crime reporting is accurate in St. Louis and Detroit, but such events are severely underreported in other places. Perhaps St. Louis and Detroit have high assault rates but few murders. Or perhaps the data used to

compute violent crimes per capita underestimate the populations of St. Louis and Detroit due to recent population growth.

A more likely problem is that there is something arbitrary in how cities are defined. City limits are political boundaries. Some cities include primarily a central urban region and exclude the outlying suburbs. Others encompass the majority of the surrounding metropolitan area. This could make a huge difference when tallying violent crime rates. For complex reasons, crime rates within many US cities tend to be high in the historical urban core of a city and lower in the suburbs.

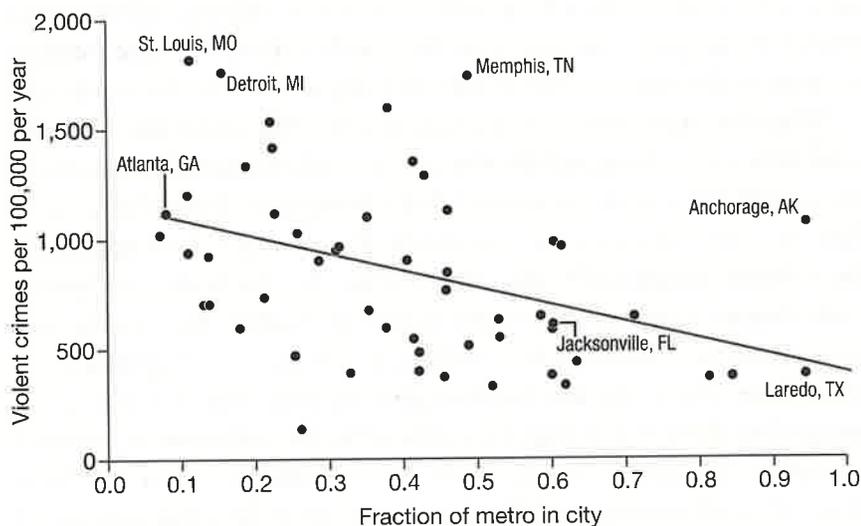
Why does this matter? It matters because the crime rate in a city will depend on how tightly the city boundary encircles the urban core. And because the locations of city boundaries depend on a city's history and politics, we see a great deal of variation in how tightly the boundaries circumscribe the city center. At this point we have a hypothesis—that city boundaries have a substantial effect on the violent crime rate—but no hard evidence. We have some grounds for skepticism about top ten most-dangerous lists,* but if we hope to argue that the way city limits are defined impacts the results, we need to collect additional data and test this hypothesis directly. Violent crime data are readily available, as are population data. But how do we control for whether the city boundaries include the suburbs?

The US government compiles a list of *metropolitan statistical areas* and assembles statistical and demographic data about each. Each metropolitan area comprises a core city or multiple core cities, surrounded by outlying suburbs. If differences in violent crime rates are influenced by differences in the way city boundaries are drawn, we would expect cities that are small compared to the surrounding metropolitan areas to have higher crime rates, on average, than cities that are big compared to their surrounding metropolitan areas.

In the scatter plot below, each dot represents one city. On the verti-

* The FBI itself urges this sort of caution in interpreting their own data: "Each year when *Crime in the United States* is published, some entities use the figures to compile rankings of cities and counties. These rough rankings provide no insight into the numerous variables that mold crime in a particular town, city, county, state, tribal area, or region. Consequently, they lead to simplistic and/or incomplete analyses that often create misleading perceptions adversely affecting communities and their residents. Valid assessments are possible only with careful study and analysis of the range of unique conditions affecting each local law enforcement jurisdiction. The data user is, therefore, cautioned against comparing statistical data of individual reporting units from cities, metropolitan areas, states, or colleges or universities solely on the basis of their population coverage or student enrollment."

cal axis, we show the violent crime rate (measured as the number of violent crimes reported per 100,000 people per year). On the horizontal axis we show the fraction of the metro's population that resides within its major city.* This gives us a measure of how narrow or expansive the city's boundaries are.



As we suspected, the higher the fraction of the metropolitan area that lies within the city limits, the lower the violent crime rate tends to be. Cities that have narrow boundaries tend to have higher crime rates, and cities with expansive boundaries tend to have lower crime rates. If we fit a line representing crime rates through the points, that line slopes downward. Statistical tests show that this slope is likely to be meaningful rather than merely the result of chance.[†] Thus there is a correlation between the violent crime rate in a city and the way the city boundaries are drawn. The overall amount of crime in a metropolitan area influences whether a city appears dangerous or safe—but so does the way in which the city's boundaries are drawn. People are making apples-to-oranges comparisons when comparing cities such as

* We look only at cities with populations of at least four hundred thousand people. We omit major metropolitan areas that contain two or more cities of this size, because the ratio of city size to metro size will be an underestimate of the fraction of the population living in an urban core.

† The slope of a linear regression through these points is significantly different from zero at the $p = 0.005$ level. $R^2 = 0.17$.

St. Louis or Detroit, which include only the urban core, with cities such as Anchorage or Laredo, which include the suburbs as well.

This example of violent crime rates serves to illustrate a more general principle: Ranked lists are meaningful only if the entities being compared are directly comparable.

3. IF IT SEEMS TOO GOOD OR TOO BAD TO BE TRUE . . .

Early in 2017, the Trump administration instituted a set of policies restricting travel and immigration to the United States. Trump's policies influenced many aspects of American life, including higher education. In March 2017, NBC News sent out a message on Twitter about the consequences of these policy changes: "International student applications are down nearly 40 percent, survey shows."

The tweet linked to a news story and was widely shared on Twitter. But the claim it put forth seems implausible. Sure, Trump's travel ban and associated changes to US immigration policy were unlikely to make the US seem more welcoming to international students. But a catastrophic 40 percent drop in applications struck us as too large to be real. Not only is the size of the effect massive, its timing is suspect. Applications to many US universities would have been due in December or January, before Trump took office. We were skeptical.

Our skepticism follows from a general principle for spotting bullshit: *If a claim seems too good—or too bad—to be true, it probably is.* We all apply this rule of thumb on a regular basis in our daily lives. How many people actually think they've won a free vacation when they get the robocall from a telephone solicitor?

So how do we figure out whether the 40 percent statistic that NBC provides is actually bullshit? Simple: *Dig to the source.* Don't trust the tweet. In a world dominated by social media where any information we receive has already been rewritten, rearranged, and reprocessed, it's important to cultivate the habit of digging to the source.

The NBC tweet provided a link to its source, an *NBC Nightly News* story titled "Survey Finds Foreign Students Aren't Applying to American Colleges." This story explains that foreign applications are down at a number of schools, and attributes this decline to Trump's travel ban and anti-immigration rhetoric:

Applications from international students from countries such as China, India and in particular, the Middle East, are down this year at nearly 40 percent of schools that answered a recent survey by the American Association of Collegiate Registrars and Admissions Officers.

Educators, recruiters and school officials report that the perception of America has changed for international students, and it just doesn't seem to be as welcoming a place anymore. Officials point to the Trump administration's rhetoric surrounding immigration and the issuing of a travel ban as having an effect.

But hold on! This article is telling a different story than the tweet suggested. The tweet claimed applications were down by 40 percent. The story suggests that applications are down *at 40 percent of schools*. That's a different matter altogether. Applications could be down just a small percentage at those schools, for example, resulting in a very small overall drop in foreign applications. Already we've found a discrepancy between the tweet and the news story it advertises.

But which is right? The tweet or the news story? To figure that out, we have to keep digging. The news story cites a bulletin from the American Association of Collegiate Registrars and Admissions Officers; a bit of searching leads us to this report, posted eleven days before the NBC story, which provides a critical detail.* Yes, international applications decreased at 39 percent of universities—but they increased at 35 percent of universities. Taken together this isn't news, it's statistical noise. Given the information presented in the article, there is no meaningful indication of a "Trump effect" on international applications. Most likely, these numbers merely reflect chance fluctuations in the number of applications to different US schools.

So how did all of this happen? There seem to have been lapses at multiple levels. First, the NBC article is misleading because it describes only the fraction of schools with declining applications, and fails to mention that a comparable fraction of schools received an in-

* Even the original American Association of Collegiate Registrars and Admissions Officers article suffers from a misleading headline, "International Student Applications Decline, Concerns about Visa and US Political Climate Rise." The data reported do not demonstrate a net decrease in US international student applications, nor is the fraction of institutions with increased enrollment substantially different from that with decreased enrollment.

creasing number of applications. We can imagine how that would have come about. A large-scale survey reveals no evidence of systematic change in international applications in response to Trump policies; that's hardly an exciting news story. Either to liven up the story, or simply because writers or editors lack quantitative sophistication, they highlight the decline in applications at 39 percent of schools and ignore or downplay the increase in applications at 35 percent of schools. Here is an example of how a statement can be true and still qualify as bullshit. It is true that applications declined at 39 percent of schools. But the absence of context is bound to mislead the reader.

The new story was then presumably misinterpreted by whoever runs the social media feed for NBC, and a decline *at* 40 percent of schools was transposed into a decline *by* 40 percent. This improbably large effect is where the "if it sounds too good or bad to be true . . ." rule of thumb comes into play. We find that this third rule is particularly good for spotting bullshit that spreads across social media. In a social media environment, the posts that are spread most widely are often those that shock, spark a sense of wonder, or inflame feelings of righteous indignation: namely, those that make the most extreme claims. And the most extreme claims are often too good or too bad to be true.

4. THINK IN ORDERS OF MAGNITUDE •

Think back to philosopher Harry Frankfurt's distinction between bullshit and lies. Lies are designed to lead away from the truth; bullshit is produced with a gross indifference to the truth. This definition gives us a considerable advantage when trying to spot bullshit. Well-crafted lies will be plausible, whereas a lot of bullshit will be ridiculous even on the surface. When people use bullshit numbers to support their arguments, they are often so far off that we can spot the bullshit by intuition and refute it without much research.

The National Geographic Society sent out a mailer warning that plastic waste is polluting our oceans. "9 Billion Tons of Plastic Waste End Up in the Ocean Every Year," the headline proclaimed. That sounds dreadful, but pause and think for a moment. There are fewer than eight billion people on the planet. Is it really possible that each person puts an average of one ton of plastic waste into the ocean each

year? That seems unlikely. In fact, the total production of plastic *throughout all history* is only about eight billion tons—and not all of that ends up in the oceans. Clearly the figure of nine billion tons per year is in error. What is the correct figure? *National Geographic* themselves recently reported that nine million tons of plastic waste go into the ocean each year. Plastic pollution of our oceans is surely an ecological disaster in the making—but inflating the magnitude a thousandfold doesn't help anyone. It merely undermines the credibility of pro-environmental sources. Not that there is any reason to suspect that the mistake was intentional; we suspect that in the production of the mailer, someone accidentally typed “billions” instead of “millions.”

Because the nine billion tons of plastic waste per year is readily compared with the eight billion people on Earth, this mistake is pretty easy to catch without doing any mental math. Often, however, one needs to do a few simple mental calculations to check a numerical claim. For example, suppose a friend claims there are more than 121,000 men in the UK named John Smith. Does that sound right? We think that the key to working out these kinds of questions quickly, without even pencil and paper, is to break down a number into components that you can estimate. The estimates can be very loose; it is usually good enough to estimate to the nearest power of ten (sometimes called an “order of magnitude”). Here we might ask, “How many people are there in the UK? What fraction of those are named John? What fraction of the UK Johns have the surname Smith?”

So how many people are in the UK? About 1 million? 10 million? 100 million? 1 billion? Most of us know that 100 million is the best estimate from among these (the true value in 2018 was about two-thirds of that, 67 million).

How many of those people have the first name John? One in ten? Well, since few women are named John, that would require that about one in five men be named John. Not even close. (Remarkably, until around 1800 one in five men in England *was* named John—but that isn't true today.) One in a thousand? Clearly the name John is a lot more common than that. One in a hundred sounds about right.

How many people in the UK have the last name Smith? Again, one in ten seems too many and one in a thousand seems too few, so one in a hundred makes for a good guess.

So how many John Smiths do we expect in the UK? To make the calculation easy we will assume that people with the surname Smith are just as likely as anyone else to have the given name John, though in practice this is probably not quite right. But we are only approximating; it should be fine to make this assumption. Therefore we have roughly 100 million people in the UK, of whom one in a hundred are named John. This gives us 1 million Johns. Of these, we estimate that one in a hundred is named Smith, giving us an estimate of ten thousand John Smiths in the UK.

This estimate turns out to be pretty good. In practice, there are about 4,700 people named John Smith living in the UK today. If we'd used the actual UK population of 67 million, we'd have had an even closer estimate of 6,700. But either way, we can see that our friend's claim of 121,000 John Smiths is off by a factor of ten or more.

This process of making back-of-the-envelope approximations is known as Fermi estimation, after the physicist Enrico Fermi, who was famous for estimating the size of an atomic blast using these simple methods.* For spotting bullshit on the fly, we suggest that making these approximations in powers of ten is often good enough. Freeing yourself up to be within a power of ten in your estimates encourages you to think through a problem quickly, using information you already know, instead of getting hung up on calculating the number of seconds in a fortnight (1,209,600 seconds) or using a search engine to find how many gallons of water the average New Yorker uses in a day (115 gallons). Even if you're off by 50 percent here and there, your final estimate is very likely to be within tenfold of the true value, which is often sufficient to spot bullshit. Of course, if your estimate leads you to believe a figure is nonsense, and you want to be sure, you can always look up the true numbers or make a more accurate estimate using pen and paper.

At a May 2018 hearing of the US House Committee on Science, Space and Technology, Representative Mo Brooks (R-Ala.) speculated that perhaps rising sea levels could be attributed to rocks falling into the ocean. For example, he asked his constituents to consider the continuously eroding White Cliffs of Dover. These

* For a delightful introductory course on Fermi estimation, see Lawrence Weinstein and John A. Adam, *Guesstimation* (Princeton, N.J.: Princeton University Press, 2008).

have to be filling up the ocean over time, and all the water that they displace must be going somewhere. It is comforting that like Aesop's crow,* Representative Brooks understands the consequences of putting rocks in water. But this is an entirely inadequate explanation that belies Representative Brooks's ability to grasp how vast the oceans are.

The oceans take up roughly two-thirds of the earth's surface and run to an average depth of about two miles. That is an enormous amount of water, spread over an almost unimaginably large area. Given that, how much of an effect could collapsing cliffs have?

We can work this out in straightforward fashion. Imagine that tomorrow the entire White Cliffs of Dover, and all the land for a kilometer inland, fell into the sea in a cataclysmic collapse and displaced an equivalent volume of seawater. Setting aside the gargantuan tsunami that would savage Calais and the northern coast of France, what would happen to sea levels globally?

Would the rising waters flood low-lying areas of coastal cities? Hardly. We can see this with a simple calculation. The White Cliffs are a bit over 10 kilometers in length and roughly 100 meters high. So in our imagined collapse, we have $10 \text{ km} \times 1 \text{ km} \times 100 \text{ m} = 1 \text{ billion cubic meters}$. Wow! It would take the Crayola company about two million years to package that much chalk.†

But the Earth's oceans are all connected, so this billion cubic meters of land would have to raise the water level of all the world's oceans. The surface area of these oceans is approximately 360 million square kilometers, or 360 trillion square meters. (We care about the surface area of the oceans, not the volume, because it is the surface that is raised.) So the 1 billion cubic meters of water displacement would be spread across 360 trillion square meters. The ensuing sea level rise would be $1,000,000,000 \text{ m}^3 / 360,000,000,000,000 \text{ m}^2 = 0.000003 \text{ m}$.

* In one of Aesop's fables, a thirsty crow adds stones to a pitcher of water until the water level rises high enough for the bird to enjoy a drink. Unlike many of Aesop's other fables, this one may be a natural history observation rather than a moral lesson. A recent scientific study showed that crows indeed have good intuitions about what happens to the water level in a vessel when you submerge heavy objects in it.

† Crayola produces about 110 million sticks of chalk every year. If each stick weighs 10 grams, that's roughly 100 million kilograms of chalk. Chalk weighs about 2,000 kilograms per square meter, so Crayola is producing very roughly 500 cubic meters of chalk each year. At 500 cubic meters a year, it would take nearly two million years to produce the full 1,000,000,000 cubic meters of chalk that we are imagining might fall into the ocean in a White Cliffs catastrophe.

In other words, we are looking at a 3-micrometer (μm) rise in sea level. By comparison, a human hair is about 100 μm in diameter. So if the White Cliffs of Dover and thirty other comparably sized stretches of shoreline fall into the sea tomorrow, the world's oceans will rise by a hair. Literally.

In practice, when White Cliffs fall into the sea they do not do so a kilometer at a time. Rather, the average rate of erosion is about 1 cm per year. That means that each year sediment falling from the White Cliffs displaces about $10 \text{ km} \times 1 \text{ km} \times 1 \text{ cm} = 100,000$ cubic meters. Spread across the whole of the world's oceans, this corresponds to a sea level rise of $100,000 \text{ m}^3 / 360,000,000,000,000 \text{ m}^2 = 0.0000000003 \text{ m}$. This is 3 angstroms—which happens to be about the size of a water molecule. So, very approximately, the annual sea level rise due to erosion from the White Cliffs is about the height of a single water molecule everywhere across the world's oceans.

Fermi estimation is useful for more than scientific problems. The same approach provides a powerful way to think about social issues. By way of example, in late 2016 the Fox News Channel's *Fox and Friends* program ran a story about food stamp fraud in the US, and framed this as a possible reason to eliminate the food stamp program (now known as SNAP). The story claimed that food stamp fraud had reached an all-time high, and that USDA statistics revealed a loss of \$70 million to fraud in 2016.

Seventy million dollars! Wow, that's a lot of money. Sounds like a disastrously run program, right? Maybe even worth canceling, given that we're wasting government funds on scammers "including a state lawmaker and even a millionaire"?

Well, this is where Fermi estimation comes in handy. First of all, you may not know exactly how expansive the food stamp program is, but you could probably estimate that about 10 percent of Americans are on food stamps—or at least that it's closer to 10 percent than to 1 percent or 100 percent. (It's actually around 15 percent.) Second, you're probably aware there are about 300 million people in the country. So around 30 million people are on food stamps. The actual number is about 45 million, but our estimate is plenty close for Fermi purposes.

If you are unfamiliar with the US food stamp program, you may not have a good idea of the average benefit paid to recipients annually.

Still you could probably guess that it's closer to \$1,000 than to \$100 or \$10,000. (In fact, it's about \$1,500.)

At this point, you've got enough information to see what's wrong with Fox's argument. Using your Fermi estimates, the US invests approximately $30,000,000 \text{ people} \times \$1,000/\text{person} = \$30,000,000,000$ —thirty billion dollars—in its food stamp program. That means that the fraction $\$70,000,000 / \$30,000,000,000 = 0.0023$, or around 0.2 percent, is lost to fraud. Using the actual annual expenditures, the fraction turns out to be less than 0.1 percent, but your Fermi estimate is plenty good to see what's going on. If there is any substantive inefficiency in the food stamp program, it certainly isn't fraud. These figures would be the envy of any retailer; retail losses to “shrinkage” (fraud, shoplifting, employee theft, etc.) usually run around 1 to 3 percent of sales.

Based on the *Fox and Friends* story, we do not know the average amount stolen by each fraudulent user, but even if SNAP frauds are receiving no more than the average legitimate recipient, fraudulent actors represent a tiny fraction of SNAP benefit recipients. It would take an exceptionally punitive mindset to starve nine hundred and ninety-nine people in an effort to safeguard ourselves against one petty crook.

There is an entertaining postscript to this story. It turns out that *Fox and Friends* was incorrect with their figure of \$70 million lost to fraud, and the US Department of Agriculture demanded a correction shortly after the story ran. The funny part is that Fox's number was not too high, it was too low. Over the period from 2009 to 2011, for example, the USDA estimated that losses to just one single form of fraud, in which food stamp recipients sell their benefits to retailers for cash, account for about \$900 million annually. This loss rate is within the normal range for retail.

If you are going to make up a number out of whole cloth, be sure to make up one that actually supports your argument.

5. AVOID CONFIRMATION BIAS

Extrême claims do well on social media; so do posts that reaffirm things about the world that we already believe to be true. This brings

The implication of our friend's tweet was that this study had found large and systematic differences in how letter writers describe men and women as candidates. From the image that he shared, it appears that writers use words associated with exceptionalism and research ability when describing men, and words associated with diligence, teamwork, and teaching when describing women. If true, this could have huge consequences for the hiring process.

This tweet clearly struck a nerve. It was shared—retweeted—more than two thousand times, in part because it captures a truth that many people experience daily. There is indeed gender bias in academia, advantaging men over women in myriad ways. Given that the tweet confirms our preexisting notions about gender bias, why were we suspicious about it? First, we have worked hard to train ourselves to avoid confirmation bias. We aim to be particularly scrupulous in questioning claims that, like this one, reflect our preexisting notions about how the world works. Second, this claim falls afoul of one of our previous rules: If a claim seems too bad to be true, it probably is. The pattern shown in the tweet is remarkably strong. Virtually all of the words on the male side of the graphic refer to excellence or exceptionalism, while all of the words on the female side refer instead to some aspect of teamwork or diligence. In our experience, patterns based on human behavior tend to be noisy. We might expect to see a general tendency toward one or the other types of description for each gender, but we'd expect some crossover.

To check up on this claim, we traced back to the source, the original research paper. While the tweet suggests that there are shocking differences between how men and women are described, the conclusions of the paper suggest otherwise:

Overall, the results of the current study revealed more similarity in the letters written for male and female job candidates than differences. Male and female candidates had similar levels of qualifications and this was reflected in their letters of recommendation. Letters written for women included language that was just as positive and placed equivalent emphasis on ability, achievement, and research.

So what is going on in the figure? Why does it give such a different impression? The answer is simple: The figure illustrates the *hypothesis* of the study, not its results.

Our friend got this confused. The words labeled as male-associated in the tweet are words that the researchers selected as “standout words” (“exceptional,” “wonderful”), “research words” (“data,” “publication”), and “ability words” (“talented,” “skilled”). The words labeled as female-associated are those selected as “grindstone words” (“dependable,” “hardworking”) and “teaching words” (“communicate,” “instruct”). The researchers had hypothesized that letter writers would use more standout words, research words, and ability words when writing letters for men, and more grindstone words and teaching words when writing for women. Instead, they found that the frequencies of ability, research, teaching, and grindstone words were comparable for both genders. Only the frequency of standout words differed. So there may be some gender differences in the text of recommendation letters, but not of the kind or at the magnitude suggested by the tweet.

6. CONSIDER MULTIPLE HYPOTHESES

In this chapter we have mainly looked how you can spot bullshit in the form of incorrect facts. But bullshit also arises in the form of incorrect explanations for true statements. The key thing to realize is that just because someone has *an* explanation for some phenomenon doesn't mean that it is *the* explanation for that phenomenon.

In May 2018, TV personality Roseanne Barr posted a racist message to Twitter. Outrage ensued and Barr apologized, blaming her actions on the sleeping medication Ambien. But it was too late; the Disney-owned ABC network canceled her sitcom despite a record-setting comeback.

Whatever one thinks about Roseanne, racism, Twitter, or Ambien, what happened next was interesting. The leading newswire service Reuters reported on Twitter that “JUST IN: Walt Disney shares down 2.5 percent after ABC cancels ‘Roseanne’ show.” Reuters was correct that Disney stock dropped 2.5 percent that day—but the headline here implies that the *Roseanne* cancellation was somehow re-

sponsible for the decline in share price. It couldn't have been: The 2.5 percent drop occurred before, not after, the *Roseanne* announcement. Indeed, the stock market as a whole had plunged dramatically that morning. Disney's stock had actually ended its 2.5 percent slide prior to the *Roseanne* announcement in the early afternoon.

This is a powerful example of "if it seems too bad to be true, it probably is." Disney is an enormous conglomerate. *Roseanne* is a single sitcom. Disney generated about \$55 billion in revenue in 2017. Season ten of *Roseanne* generated about \$45 million in revenue in 2018. How could the loss of a series that generated 0.1 percent of Disney's revenue drive a 2.5 percent decline in stock price? It doesn't pass a basic plausibility check.

The problem with the Reuters tweet is that when you have a phenomenon of interest (a 2.5 percent slide in Disney's stock price) and a possible explanation for that phenomenon (*Roseanne* was canceled), your story can seem compelling. *Roseanne*'s racist tweet was by far the most societally salient event associated with Disney's holdings at the time—television pundits, newspaper columnists, and social media posters alike were up in arms about what she wrote and about the consequences she faced for doing so. But this does not mean that this event was the most plausible explanation for the trading results.

The real point here is that in the vast majority of cases, there will be many possible explanations for any given pattern or trend, and simply being consistent with the data does not make an explanation correct. Sometimes a proposed explanation may be correct but only a small part of the causal contribution; other times the proposed explanation may be totally wrong, unconnected to the real explanation.

In order to avoid falling for these seemingly plausible but incorrect explanations, consider as many possible explanations as you can for any trend or pattern or event that you seek to understand.

SPOTTING BULLSHIT ONLINE

In chapter 2, we described how social media has changed the way that news—real or fake—spreads. When we decide what is worth sharing

on Facebook, Twitter, or some other platform, we are taking on the gatekeeping role that professional editors once played. We are not only fooled ourselves by online misinformation; we are readily enlisted as vehicles to help spread it. That makes it particularly important for us to get good at spotting bullshit when it appears on the Internet or in our social media feeds. We conclude this chapter by summarizing the suggestions we've presented throughout the book for spotting online misinformation. Some of them may sound simple. But simple as they may be, the reminder is useful. We have also found that we need to continually work to cultivate our own habits of mind, including the techniques we list below.

1. Corroborate and triangulate. If you come across a surprising claim or dramatic news report from an unknown source, use a search engine to see if you can find the same claims from other sources. If not, be very suspicious. Even when one news outlet has a big scoop, other papers quickly report on the fact that the first outlet broke the story. Be sure that those reporting on the story include reliable sources. Disinformation campaigns may plant multiple versions of the same false tale in unreliable outlets.

2. Pay attention to where information comes from. If you find a piece of candy lying in the street, you are not going to eat it or share it with your friends. Unsourced information is the same. All too often, someone we don't know posts a factoid or statistic or data graphic on social media without listing its source—and we share it anyway.

3. Dig back to the origin of the story. This takes time and effort, but if you want to avoid spreading misinformation, it is effort well spent. Don't simply read a headline or tweet; read the full news story. If the news story is from an outlet that tends to sensationalize, don't stop there. Dig back to the primary article or report that the story is talking about. Or dig really deep and take a look at the data yourself.

4. Use reverse image lookup. Several search engines provide a reverse image lookup service in which you upload a picture or a few frames from a video, and the search engine tells

you where on the Web that picture or video can be found.* This is one of the more underutilized tools on the Web for fact-checking. If you are suspicious of a Twitter or Facebook account, check to see if the profile photo comes from a stock photo website.

5. Be aware of deepfakes and other synthetic media. A random stranger on the Internet could be anybody, anywhere. But while we've learned to distrust user names by themselves, we're still susceptible to people's pictures. In the past, a person's photo was pretty good proof that they existed. No longer. So-called deepfake technology makes it possible to generate photorealistic images of people who don't exist. For now, one can still spot them with a bit of practice. Learn how at our website, <http://whichfaceisreal.com>. It won't be long until these fakes are much more difficult to detect, so it's good to keep in mind that even if someone appears in a "photograph," he or she still might not be a real person.

6. Take advantage of fact-checking organizations. If you come across a wild story online, confirm it by visiting a fact-checking website such as Snopes.com, PolitiFact.com, or FactCheck.org. If the story is not yet documented at these websites, ask them to fact-check it. They learn what stories need to be verified or debunked from users like you.

7. Make sure you know whom you are dealing with. Like other Internet fraudsters, fake news creators try all kinds of tricks to make the information they are providing seem more legitimate than it is. A fake news story might be designed to look as if it is just one of hundreds of stories from a large newspaper or television news station—but if you were to dig deeper, you would find that no such paper or station exists. Alternatively, a fake news pusher might send a social media post with a link that says something like "view the story on cnn.com," but the link would actually direct you to a different

* TinEye is a standalone reverse image search engine: <https://tineye.com/how>.

Google provides instructions for using their reverse image search here: <https://support.google.com/websearch/answer/1325808>.

Bing's image match service is detailed here: <https://blogs.bing.com/search/2014/03/13/find-it-faster-with-image-match/>.

Web domain with pages made to look like CNN. Be aware of the web addresses you are visiting. Scammers often pick domain names that are easy to misread. Similar as they look, abc.com.co is not abc.com; faceb000k.com is not facebook.com. There are thousands of these kinds of websites that try to look legitimate. Sometimes fake news sites run ads that look like they are coming from reputable outlets but instead are traps that send you to scammers' sites.

8. Consider a website's track record. How do you know if a website is reliable? Try to find out if the site has been known to create and push fake news sources. Wikipedia often provides an overview of media outlets; this can be a good place to start. No one gets the facts right all the time, so see if the website issues corrections. Is the site reflective about the challenges it faces in getting at the truth?

9. Be aware of the illusory truth effect. The more often you see something, the more likely you will be to believe it. We take this very seriously when studying fake news and conspiracy content. We know that it can be disorienting to mill through fake news stories, so be cautious. Watch out for a tendency to believe something because you keep seeing it.

10. Reduce your information intake. Take a break; be bored a few times a day and revel in "missing out" instead of being anxious about what you missed. This will enhance your ability to process information with skepticism when you are online.

Most important: When you are using social media, remember the mantra "think more, share less." The volume of information on social media, and the speed at which it allows us to interact, can be addictive. But as responsible citizens, we need to keep our information environments as clean as possible. Over the past half century people have learned not to litter the sides of the interstates. We need to do the same on the information superhighway. Online, we need to stop throwing our garbage out the car window and driving away into the anonymous night.