

Refuting Bullshit

WE BEGAN THIS BOOK BY DEFINING BULLSHIT, DISCUSSING ITS origins deep in our evolutionary history, and explaining how and why it proliferates so readily in today's digital environments. We looked at the various forms that bullshit—particularly quantitative bullshit—can take, and how to detect them. The book so far could be called *Spotting Bullshit*. But we chose to call it *Calling Bullshit*, because a solution to the ongoing bullshit epidemic is going to require more than just an ability to see it for what it is. We need to shine a light on bullshit where it occurs, and demand better from those who promulgate it.

We define *calling bullshit* as follows:

Calling bullshit is a performative utterance in which one repudiates something objectionable. The scope of targets is broader than bullshit alone. You can call bullshit on bullshit, but you can also call bullshit on lies, treachery, trickery, or injustice.

This definition draws upon the concept of a performative utterance, as discussed in the philosophy of language. When we think about the purpose of language, we usually think about how we use it to make statements about ourselves or about the world. "I feel sad." "The next bus doesn't come until seven-thirty." "Many years later, as he faced the firing squad, Colonel Aureliano Buendía was to remember that distant afternoon when his father took him to discover ice."^{*}

* So begins Gabriel García Márquez's *One Hundred Years of Solitude*.

Words are also useful for issuing commands: “Stop!” “Flight attendants, prepare for landing.” “Honour thy father and thy mother: that thy days may be long upon the land which the Lord thy God giveth thee.”

In the wryly titled book *How to Do Things with Words*, philosopher J. L. Austin noted that there is yet a third class of things that we do with speech. There are utterances that, when we make them in the appropriate circumstances, are better viewed as actions than as expressions of a proposition. These are known as performative utterances. “I dub thee Knight”; “I christen this ship the HMS *Beagle*”; “I do [take this man as my lawfully wedded husband]”; “I do solemnly swear that I will support and defend the Constitution of the United States against all enemies, foreign and domestic.” In each case, the speaker is not merely reporting on her action, she is acting by means of her speech. Austin calls these sentences performative utterances, because one performs an action by uttering the expression.

Performative utterances are statements rather than questions. The subject is usually “I,” and they are in the present tense rather than past or future tenses: “I resign” instead of “I resigned” or “I will resign.” You could summarize all of this by saying that performative utterances tend to be expressed in the first-person singular, indicative active present tense—if you’re into that sort of thing. In addition to grammatical cues, the English language even has a somewhat archaic word, “hereby,” that can be used to flag a performative utterance if it is not obvious from context. We don’t claim the passenger seat by shouting “I hereby call shotgun,” but “hereby” remains common in legal language, where it serves to indicate that a legal document represents an official action or commitment: “I hereby accept the agreement,” “I hereby renounce all claim upon the estate,” “I hereby declare that the details furnished above are true and correct to the best of my knowledge and belief.”

Calling bullshit is itself a performative utterance—and this observation is important for understanding what it means to call bullshit upon some claim. When I call bullshit, I am not merely reporting that I am skeptical of something you said. Rather, I am explicitly and often publicly pronouncing my disbelief. Why does this matter? Performative utterances are not idle talk. They are powerful acts, to be

used with prudence. Calling bullshit is the same. Don't call bullshit carelessly—but if you can, call bullshit when necessary.

We are convinced that the world would be a better place if there were less bullshit to deal with. As the legendary journalist and political commentator Walter Lippmann noted a century ago, "There can be no liberty for a community which lacks the means by which to detect lies." Calling bullshit is crucial to the healthy functioning of a social group, be it a circle of friends, a community of academics, or the citizenry of a nation. Any group adopts wrong ideas at times, and these ideas require forceful public repudiation. But if you want to call bullshit, it is important to do so responsibly, appropriately, and respectfully. This is not an oxymoron. We the authors do this for each other on a daily basis. We understand that the proper target of calling bullshit is an idea, not a person. We realize that we will sometimes be on the producing end of bullshit. We've learned to accept and acknowledge our mistakes with a modicum of grace when bullshit is called on us.

Spotting bullshit is a private activity. Calling bullshit is a public one. If you can spot bullshit, you can keep yourself safe from its effects. If you can call bullshit, you can protect your entire community. Of course not everyone will feel comfortable calling bullshit—and that's fine. Moreover, we recognize that it might be less acceptable to do this in some cultures. There are many ways to help reduce the bullshit density of our society without sticking your neck out too far. You can learn to spot bullshit and avoid being misled yourself. You can learn to stop producing bullshit. You can learn to avoid sharing it. But we've already talked about how to do these things. For those of you who want to take the leap and call bullshit, we'll show you how to do so both effectively and appropriately.

Carelessly calling bullshit is a quick way to make enemies of strangers and strangers of friends. Because calling bullshit is a performative utterance, it is particularly important to be correct when you do so. People despise hypocrites, and being full of shit when you call bullshit is about as hypocritical as one can get. It's worse if you've called it aggressively and judgmentally. There's a thin line between being a tough-minded skeptic and a domineering jerk. We want to make sure you don't end up on the wrong side of the line.

Why did we wait until the final chapter to talk about *calling* bullshit? To spot bullshit you need to develop all of the skills and habits of mind that we've described over the past ten chapters. You need to be aware of the traps and sources of confusion, the potential misuses of numbers and data visualization and big data, the ways in which bullshit slips into not only the popular media but the scientific literature as well. To call bullshit seems comparatively easy. You simply have to open your mouth or type a few characters to say "bullshit!" But simply uttering the words is not enough. In order to call bullshit effectively, you need to know how to construct a successful refutation. The right approach will depend not only on the type of bullshit you aim to refute, but also on the audience you wish to convince. Various approaches may be best for convincing your child, your congressional representative, a stranger on an airplane, or a scientist with a background in statistics.

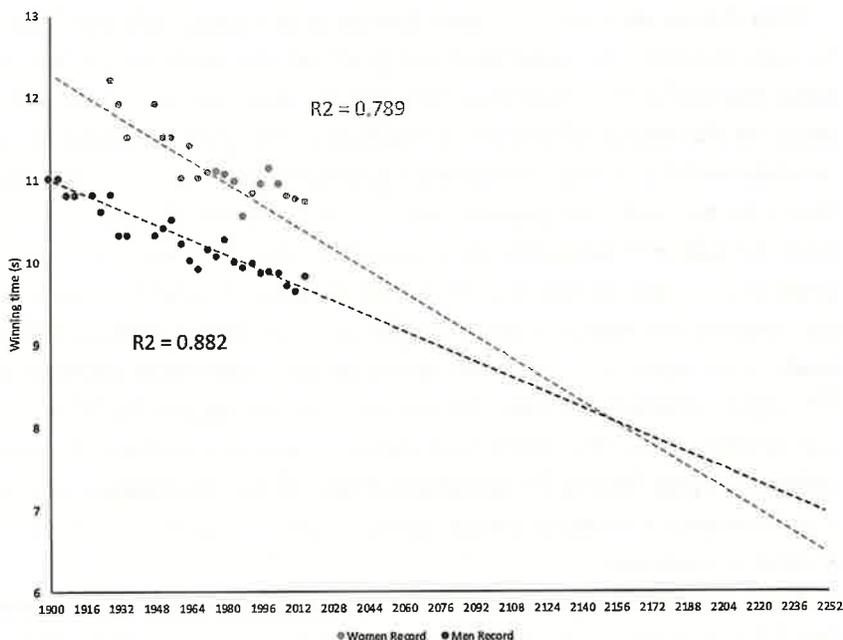
Now we turn to methods for refuting bullshit. Many of these should be familiar—we have been demonstrating them throughout the book.

USE REDUCTIO AD ABSURDUM

In the Summer Olympics of 2004, Yuliya Nesterenko won the women's hundred-meter dash in 10.93 seconds. This was not an Olympic record, but it was more than two seconds faster than the women who had run the event seventy years earlier.

Inspired by this substantial improvement over a relatively short amount of time, researchers published a short news piece in *Nature*. Comparing times for men and women, they found that over the past few decades women have been closing the gap with male sprinters. While men's times were improving, women's times were improving more rapidly. What should we expect to happen in the future? the researchers asked. Modeling the changes in winning times, the authors predicted that women will outsprint men by the 2156 Olympic Games.

It may be true that women will someday outsprint men, but this analysis does not provide a compelling argument. The authors' conclusions were based on an overly simplistic statistical model.



As shown above, the researchers fit a straight line through the times for women, and a separate straight line through the times for men. If you use this model to estimate future times, it predicts that women will outspurt men in the year 2156. In that year, the model predicts that women will finish the hundred-meter race in about 8.08 seconds and men will be shortly behind with times of about 8.10 seconds.

Of course, both women and men will continue to break records. However, there is something clearly wrong with the model. Many readers, including students in a high school class in Texas, caught this problem and wrote to *Nature* in reply. Our favorite response was a letter from a biostatistics professor, Ken Rice (emphasis added):

Sir—A. J. Tatem and colleagues calculate that women may outspurt men by the middle of the twenty-second century (*Nature* 431, 525; 200410.1038/431525a). They omit to mention, however, that (according to their analysis) a far more interesting race should occur in about 2636, when times of *less than zero seconds* will be recorded. In the intervening 600 years, the authors may wish to address the obvious challenges raised for both time-keeping and the teaching of basic statistics.

This response is both humorous and highly effective. In it, Rice employed one of our favorite refutation strategies: *reductio ad absurdum*. This strategy, which dates back at least to Aristotle, shows how your opponent's assumptions can lead to ridiculous conclusions. In his *reductio*, Rice employed the same model and methods found in the *Nature* paper. Using the same model, he extrapolated further into the future and came to the preposterous conclusion that late-millennium sprinters will run the hundred-meter dash in negative times. Clearly this can't be true, so we should be skeptical of the paper's other surprising results, such as the forecasted gender reversal in winning times.

Another lesson here is to be careful about what kind of model is employed. A model may pass all the formal statistical model-fitting tests. But if it does not account for real biology—in this case, the physical limits to how fast any organism can run—we should be careful about what we conclude.

BE MEMORABLE

Functional magnetic resonance imaging (fMRI) allows neuroscientists to explore what brain regions are involved in what sorts of cognitive tasks. Researchers have looked at what areas of the brain are most active in subjects playing videogames, having sex, listening to music, engaging in exercise, or responding to other stimuli. A typical study would compare the fMRI images of afflicted patients with those of nonafflicted controls and ask why certain parts of the brain light up differently.

This technique for measuring neuronal activity has left an enormous mark on neuroscience. But the software used to detect differences in brain activity has to make assumptions about how to assess the statistical significance of these results. A recent study showed how these assumptions can sometimes make differences look more significant than they are. While scientists do not universally agree on the size of the problem, some feel the problem is serious enough to call into question the results from several thousand papers.

Years before these statistical issues came to light, a research poster presented at a neuroscience conference used *reductio ad absurdum* to present a memorable critique of fMRI technology. The title of the poster? "Neural Correlates of Interspecies Perspective; Taking in

the Post-mortem Atlantic Salmon: An Argument for Proper Multiple Comparisons Correction.” You read correctly: a dead salmon.

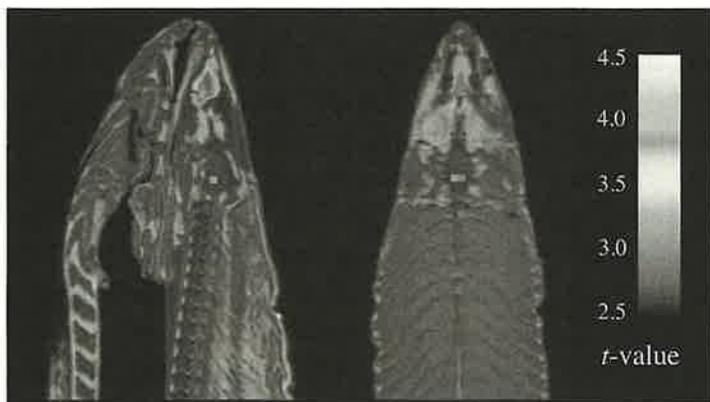
This *reductio ad absurdum* came from a deliberately silly experiment. The authors placed a dead Atlantic salmon into an fMRI machine, asked the unfortunate creature the sorts of questions that researchers typically ask human subjects, and measured the activity of the deceased fish’s brain. In a subsequent research paper, they described the study as follows:

The salmon measured approximately 18 inches long, weighed 3.8 lbs, and was not alive at the time of scanning. It is not known if the salmon was male or female, but given the post-mortem state of the subject this was not thought to be a critical variable. . . . The salmon was shown a series of photographs depicting human individuals in social situations with a specified emotional valence, either socially inclusive or socially exclusive. The salmon was asked to determine which emotion the individual in the photo must have been experiencing.

It’s brilliant—we have firsthand confirmation from one of the researchers that they actually spoke to the dead salmon and showed it pictures of people in different situations—and the results were striking. Several regions in the brain stem of the salmon showed higher activity when the fish was being asked about people’s emotions, relative to when it was at “rest.” (We can only imagine how bright the regions may have been if the salmon was asked about salmon emotions.) If these data had come out of a more straight-faced study, the researchers might have concluded that this brain stem region was involved in processing emotion. Instead, the findings highlight the risk of false positives in fMRI studies.

The authors summarized their results, in the most memorable of ways:

Either we have stumbled onto a rather amazing discovery in terms of post-mortem ichthyological cognition, or there is something a bit off with regard to our uncorrected statistical approach.



The salmon study was followed by a number of more technical and less entertaining research articles that further dissected the problem, estimated its magnitude, and proposed solutions. These projects were critical for the field to advance, but none were as effective as the original salmon study in drawing attention to the basic problem. Humor is not a requirement for *reductio ad absurdum*, but when integrated well it can be highly effective. It is memorable and spreads ideas quickly through informal conversation.

FIND COUNTEREXAMPLES

The Santa Fe Institute (SFI) is an independent scientific research center perched high above the city of Santa Fe, New Mexico, in the foothills of the Sangre de Cristo Mountains.* There, scientists from a wide range of disciplines—physics, biology, economics, sociology, computer science, mathematics, and beyond—meet to study, think, and discuss ideas over afternoon tea. It is the kind of environment where creative, big-picture thinking is encouraged and supported. But rigor is also essential, and the visitors to SFI are more than willing to point out mistakes in one's logic or holes in one's reasoning.

We heard about a great counterexample raised at a Santa Fe Institute workshop on immune systems. One physicist at the meeting

* Carl spent a decade on the external faculty at SFI, and Jevin has been a frequent visitor to the institute.

had learned a little bit about the immune system and had created a mathematical model to account for it. In his talk, he described his model and stressed its powerful implications. Not only are immune systems like ours useful for dealing with pathogens, he explained, but they are absolutely necessary. In order to survive in environments overrun with pathogens, he predicted, long-lived multicellular organisms such as ourselves must necessarily have immune systems with certain distinctive features. For example, the model suggested that long-lived organisms will need to have cells that detect the presence of viral infection, and other cells that produce a wide range of antibodies, generated randomly and then selected to proliferate should they match the invading pathogen. At first blush, the argument may have seemed reasonable. Plus, there was fancy mathematics behind it!

An immunologist in the room seemed unconvinced. It is at this point that someone will often raise their hand to question the assumptions of the mathematical model or ask for clarification of the analysis; extensive technical discussion and often disagreements on the mathematical details follow. The immunologist—with decades of experience designing and conducting experiments, reading and reviewing thousands of papers, and teaching courses on immunology to college students—took a different approach. Instead he asked a question that required no more than a basic education in biology. He raised his hand and asked the speaker: “But what about trees?”

Trees are multicellular organisms, and they are certainly long-lived. By most accounts, the bristlecones of the White Mountains are the oldest living organisms on the planet. But while trees have immune defenses, they have few if any of the characteristics that the speaker claimed were needed to keep a large organism alive for a long period of time.

The tree counterexample was devastating. There was basically no point in going any further with the talk or the questions. The speaker and the audience might as well have ended early and grabbed some of the delicious cookies and coffee that are always on hand for breaks.

Reductio ad absurdum can be fun and effective but nothing is as devastating to specious arguments as a simple counterexample. If someone claims that A implies B , find a case in which A is true but B is not. A in this case is a large, long-lived multicellular organism; B is having

an adaptive immune system. Trees can be classified as A , but do not have B ; therefore A does not imply B .*

It takes practice to find good counterexamples, and few will be as effective and crushing as the tree example. If and when you do find one, please be kind. If the claims were made in earnest with no ill intent, help your target recover. We all make mistakes. Hopefully your counterexample will be instructive and will lead to better analyses with stronger claims in the future.

PROVIDE ANALOGIES

Like many large cities, Seattle has serious traffic problems. Nearly four hundred thousand vehicles commute into the city every day, and the aging infrastructure was not designed to handle anything near that volume. Making matters worse, Seattle is one of the fastest-growing cities in the US, and as the city is surrounded on both sides by water, this growth takes the form of increasing population density rather than outward expansion.

One especially problematic region is South Lake Union near Mercer Street. This once sleepy patchwork of auto body shops and small local businesses has been replaced by soaring forests of glass and steel that house Amazon's rapidly expanding headquarters and a host of

* Nowhere is the power of counterexample as clear as in mathematics. A single counterexample provides absolute resolution to an open conjecture. Consider Fermat's last theorem (which was not really a theorem but a conjecture, because Fermat did not provide a proof). It states that there are no three distinct integers a , b , and c such that $a^n + b^n = c^n$ for integer values of n greater than 2. For centuries, mathematicians tried to prove this conjecture—and failed. Finally, after years of solitary work developing major advances in several areas of mathematics, British mathematician Andrew Wiles came up with a proof drawing on recent mathematical discoveries and running 127 pages in length. It was an epic achievement, with numerous setbacks including the discovery of a mistake in the proof that set Wiles back another two years. That's what it takes to prove Fermat's last theorem true. Had the theorem been false, the proof could have been much simpler. It could have been a mere counterexample, three numeric values for a , b , and c , and a value for n .

Indeed, this is exactly what happened when the great eighteenth-century mathematician Leonhard Euler attempted to generalize Fermat's last theorem into what is known as the sum of powers conjecture. It says that for integers a , b , c , . . . , z and any integer n , if you want numbers a^n , b^n , c^n , etc., to add to some other number z^n , you need at least n terms in the sum. For nearly two hundred years, no one was able to prove or disprove the conjecture. Then in 1966, two mathematicians used an early computer to run through a huge list of possibilities and found a counterexample: $27^5 + 84^5 + 110^5 + 133^5 = 144^5$. No further theory is needed; from that one example we know that Euler's sum of powers conjecture is false. Counterexamples can be much easier to verify than positive proofs. There are relatively few mathematicians in the world who can verify all of Wiles's proof. But you can verify the counterexample to Euler's conjecture in a minute or so with the help of an electronic calculator.

other tech firms. Simply leaving the “Mercer Mess” and getting out onto the interstate can take the better part of an hour during the worst blockages. A few years back, the city decided to invest \$74 million to improve traffic flow through this area. When the construction was finished and traffic had returned to the streets, the transportation office began measuring the travel times. The results were not what some had hoped. One local television station reported “\$74 Million Later, Mercer Mess Is 2 Seconds Faster.” The implication was obvious: yet another colossal waste of taxpayer money.

But was it really a waste? It is true that from the perspective of an individual driver, travel time did not decrease much. But think about why this happened. First of all, the number of people working in this area increased faster than almost anywhere in the country. Simply holding travel times constant in the face of this influx is a huge accomplishment. Second, the traffic patterns in a big city adjust themselves organically. If one route becomes substantially faster, traffic diverts from nearby routes to the faster route until travel times equilibrate for all routes. If you increase the capacity of one road, it may speed travel there only a little; the benefits are distributed across the entire traffic grid.

That seems to be precisely what happened. The Mercer corridor now handles thirty thousand more cars a day than it did before the improvement project, without any increase in travel time. The headline might have read “Seattle Road Improvement Project Allows 10 Million Extra Trips Per Year with No Increase in Travel Times.” In order to measure the project’s benefits, we need to consider the consequences of an improvement project across all routes in that region of the city.

Not long after we launched our Calling Bullshit course, we were asked to talk about the Mercer traffic situation on one of the local news stations. To highlight the folly of measuring travel times while ignoring trips per day, we used an analogy that resonated with Seattle-area viewers. In early 2010, the Seattle Mariners baseball team renewed the contract of pitching ace Felix Hernandez for about \$78 million—comparable to the cost of the construction work on the Mercer Mess. But the Mariners’ team batting average dropped from .258 in 2009 to .236 in 2010, and the number of home runs hit by Mariners players declined from 160 to 101. A local TV station could have run a headline

gripping that “Mariners Hitting Declines in 2010 Despite \$78 Million Investment in Felix Hernandez.” It’s a true claim. But it’s obviously a silly headline, because it couples two unrelated facts. The hitting decline has nothing to do with Hernandez. Pitchers play only every fifth game or so, they don’t bat in the American League and, most important, they are not signed for their batting prowess. Pairing Hernandez’s contract with the team’s 2010 hitting performance implies that his salary was wasted because he didn’t raise the team’s batting average. Absurd, but not so different from suggesting that the Mercer Mess construction was a waste because travel times did not appreciably decrease. Each uses an irrelevant measure to evaluate the payoff of a seventy-million-dollar investment in a Seattle institution.

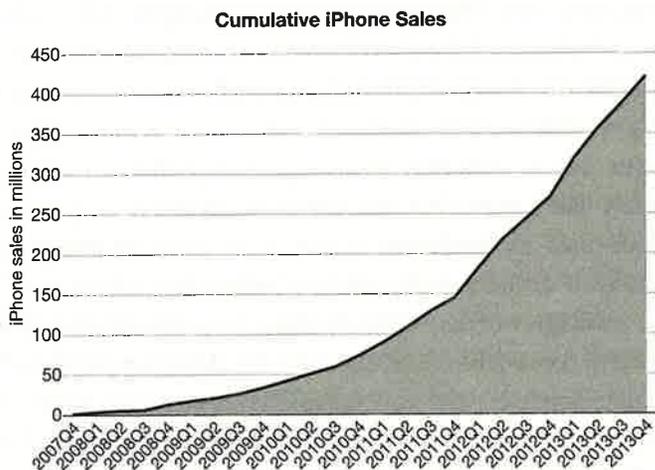
We use analogies often, because they help recontextualize claims that may seem reasonable at first glance. By drawing parallels between an unfamiliar situation and an example your audience intuitively understands, you encourage them to trust their own critical thinking abilities. For example, when we addressed the vaccine skeptic’s dismissal of the one percentage point decrease in flu, we drew an analogy to another modern safety innovation that people understand and accept: seatbelts. We showed that she could make approximately the same claim about seatbelts as she did with the flu vaccine. Parents might feel unequipped to judge the risks associated with vaccines, but seatbelts are a matter of common sense.

REDRAW FIGURES

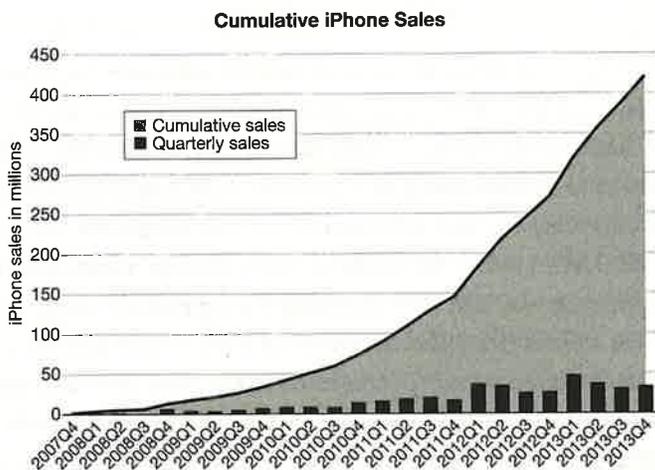
In chapter 7, we looked at ways in which accurate data can be displayed misleadingly. While one can point out some of the tricks that designers have used, the most effective way to refute the misleading impression is to redraw the graph in a more appropriate fashion.

In that chapter, we saw an example of this when we looked at the *National Review*’s tweet: the only climate change chart you need to see. By zooming way out and showing a temperature range from 0 to 110 degrees Fahrenheit, the designer of that graph hid the dramatic two-degree increase in temperature over the past few decades. With the exact same data, *The Washington Post* subsequently redrew the figure, zooming in to a more appropriate range. The temperature increase that leaps out of that chart tells a different story. 🌱

The Internet news site *Quartz* used this technique to strong effect when, in 2013, Apple CEO Tim Cook gave a presentation about iPhone sales. Below is a version of the graph that Cook showed:



This graph looks quite impressive; it seems as if Apple is taking over the world with the iPhone as cumulative sales go up and up and up. But of course cumulative sales go up—cumulative sales can't go down! What the graph hides is the fact that quarterly iPhone sales had been declining for at least the past two quarters prior to Cook's presentation.

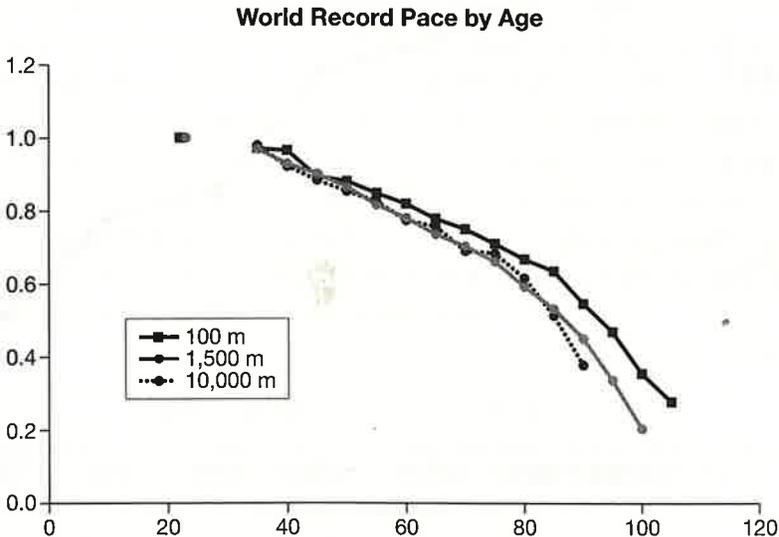


By redrawing the graph to show quarterly sales, *Quartz* was able to reveal a story that was hidden in the original line chart.

DEPLOY A NULL MODEL

Once we reach our physical peak sometime in our twenties or thirties, our performance for most physical and cognitive tasks begins to decline. Biologists call this process *senescence*.

To illustrate this lamentable fact in the evolutionary biology textbook Carl co-wrote, he plotted the world records for various field events by age. The figure below shows the average speed of world record holders in the men's 100-meter, 1,500-meter, and 10,000-meter race, with the speeds normalized so that the world record corresponds to a pace of 1.0.

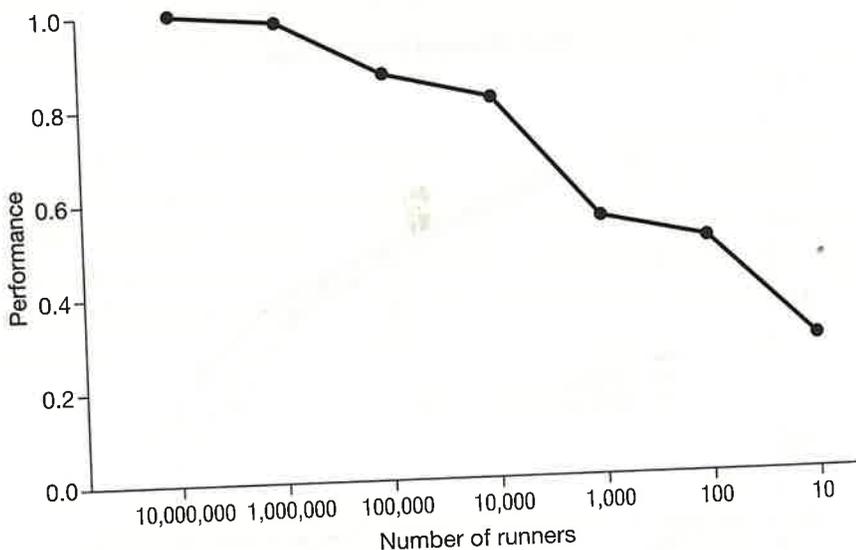


While teaching his course on evolution and medicine, Carl showed the students this graph. He explained that it was intended to illustrate that human physical performance declines with age, and asked them to take a few minutes to think about whether there were any problems with that argument. He anticipated that they would come up with a few objections. For example, these are world record times set

by the world's best athletes. The performance curve above may not be representative of what happens to the rest of us.*

What Carl didn't anticipate, however, is that one of his students would point out a problem that he had not even considered. There are many more people running competitively in their twenties and thirties than in their seventies and eighties. The more runners you sample from, the faster you expect the fastest time to be. She was completely right. The fastest runner in a sample of a million people is likely to be much faster than the faster runner in a sample of a thousand.† We might see the same decreasing trend in speed simply as a consequence of sample size, even if runners did not get slower with age. If so, Carl's graph doesn't provide very compelling evidence for senescence.

World Record Pace by Age



* Another issue with this graph is that the curve shown does not represent the performance trajectory of any single individual. Different athletes hold the records for different ages. Those who peak early may fall off faster than illustrated by the curves. Those who peak late may not have had as far to fall because they never ran at world record pace in their youth. So the curves don't really tell us about what happens to a person over time, they just tell us something about the upper bounds of human performance. Moreover, there may be "cohort effects" operating. Runners who set records in the sixty-five-and-over group trained using different techniques, diets, etc., than runners currently in their twenties. Improved training technology could have an effect on record times as well.

† If we were looking at *average speed*, the sample size wouldn't matter much. Whether we sampled a hundred, a thousand, or a million runners of a given age, we would expect the average time to be about the same. But if we are looking at the extreme outliers, the sample size matters.

To illustrate the student's point, we can create a *null model*. A null model helps us understand what we would observe in a very simple system where not much is going on. In this case we can use a computer simulation to create a pretend world where age doesn't affect running speed. Then we see if we still observe the same downward trend in physical performance simply because there are fewer runners of older ages. The graph in the previous page illustrates what we find.

Our null model generated data very much like the data that Carl presented in his textbook—without requiring any senescence. This does not mean that senescence is a myth. What it does mean is that the data Carl plotted do not provide compelling evidence of senescence, because the null model shows the same result without senescence.

That is what null models do. The point of a null model is not to accurately model the world, but rather to show that a pattern *X*, which has been interpreted as evidence of a process *Y*, could actually have arisen without *Y* occurring at all. In our example, the pattern *X* is the decreasing performance of record holders with age. The process *Y* is senescence: Humans run slower as they get older. Because we see decreasing performance even without senescence, our null model is a way of saying, "Sorry, you don't get to use those data as evidence for your theory." Notice here that the null model does not have to be an accurate description of the world to rob the original argument of its rhetorical force. It is enough to show that we would see the same pattern even without the process in question.

THE PSYCHOLOGY OF DEBUNKING

Myth is the most difficult to debunk when they are interwoven with a person's worldview and sense of cultural identity. We once had dinner with a zoo director who faced repeated challenges from PETA, People for the Ethical Treatment of Animals. He wanted to know how he could engage in a constructive dialogue and possibly even persuade them that his zoo plays a valuable role in conservation efforts. We explained how their views were entangled with their identities in a way that his were not. For example, if they were to convince him that keeping elephants in captivity was unethical, he would still be a scholar and zoo director. But if he were to persuade them that keeping elephants in zoos was justifiable, they could not retain their iden-

tities as PETA activists. These issues of identity made his task vastly more difficult. Often your task won't be difficult. Help your aunt see that she can maintain her electrolyte balance without paying eighty dollars for a water bottle with an amethyst crystal in it. Help your uncle see that he can still prefer a limited federal government without denying global warming. Whatever it takes. Find ways to decouple issues of identity from the matter you are trying to debunk.

In the previous chapter we discussed being aware of our own confirmation biases, which arise because we tend to search for, believe, and recall information that aligns with our own worldviews. Recognize that others suffer from confirmation bias as well. Once an idea is entrenched it is difficult to replace with a better supported idea, no matter how hard you try. In addition to the methods we've discussed so far, here are some time-tested tips for debunking myths:

1. Keep it simple. One advantage that falsehood has over truth is that the truth is often complicated whereas falsehoods can be crafted to be simple. Look for ways to make your story as simple as possible without distorting it. Focus on your core points and let the rest go. Scoring rhetorical points on tangential technicalities doesn't convince anyone, it just pisses people off.
2. Take it offline. Take your conversations offline if you can. No one likes to be called out in public. If you're going to call bullshit on your brother-in-law, do so on a quiet walk, not at the Thanksgiving table. If you're going to call out someone online, consider doing so in a private message rather than on someone's public Twitter feed. (Admittedly, this is complicated by a need to protect others from being deceived.) When called out publicly, most of us dig in and try to defend our previous statements rather than considering objections in good faith. Privately, people are often more open to critique.
3. Find common ground. The less antagonistic your interaction is, the more likely someone will seriously consider your ideas. One of the best ways to soften your words is to first establish a point of common ground. When talking to someone who is skeptical about vaccine safety, don't start out saying "I can't believe you're dumb enough to fall for that stupid

hippie BS.” Instead, try the common-ground approach. “Wow, it’s really hard to know how to do right by our kids. I worry about it all the time myself . . .”

4. Don’t overemphasize the myth. Familiarity increases the stickiness of myth. If a reference to the myth is necessary, precede with explicit warnings. Some research suggests that people’s beliefs in misinformation can be strengthened if you reiterate the myth before you debunk it.

5. Fill in knowledge gaps with alternative explanations. It is not enough to simply debunk a myth; you need to replace it with an alternative account. People don’t like incomplete stories in which events lack explanation. Your argument against a myth may seem convincing now, but if you don’t replace it with a new narrative, your audience may fall back on the same old misinformation in the future. Good defense lawyers know this. Instead of just explaining to a jury why their defendant is *not guilty*, they will point to other suspects or circumstances that can fill the void if their client is not the perpetrator.

Now that we have given you some pointers to spot and refute bullshit, you’re almost ready to go. But before you go out there and start calling bullshit willy-nilly, we would like to conclude with a few thoughts about how to do so in an ethical and constructive manner.

Be correct

The students who take our course on calling bullshit gain confidence in their ability to spot and refute it. We want to instill a sense of humility as well. It should go without saying, but it needs to be said: If you are going to call bullshit, be correct.

Make sure you have the facts at hand—don’t skimp on the background research—and then double-check them. Mountaineers learn to check and double-check and then check their safety gear one more time. Just as you don’t want to have your locking carabiner fly open during a forty-foot fall, you don’t want to have your facts wrong when you’re trying to drive home a point. Also, run your argument by a friend or colleague. We have been doing this for each other for fifteen years. If one of us gets excited about a new research result, he

asks the other to try to undermine it. We've saved each other from embarrassment a number of times.

Be charitable

Twitter has been compared to a nation of television viewers yelling back at their TV sets, encouraged by a faint hope that the people on the TV might hear them. This seems an apt description of a lot of public discourse today. In a combative environment like this, it is easy to demonize those who disagree with us. When we hear something that we think we know is wrong, we are tempted to invoke malice or even conspiracy as the reason for the falsehood. Consider a few alternatives:

- You might be wrong. You may think this is unlikely, but at least be mindful of the possibility. It could be that you misheard the statement or misinterpreted an argument.
- Don't impute malice when incompetence is a sufficient explanation. Most people who write foolish things on the Internet or anywhere else do not have a nefarious underlying motive. They simply don't know what they are talking about.
- Don't assume incompetence when an honest mistake can account for error. We all make honest mistakes and say stupid things at times; it doesn't mean that we are stupid or that we are incompetent.

Being charitable is about keeping your friends, but it is also about focusing your refutation on the argument itself. To call bullshit in a civil way, attack the argument, rather than the person. Your neighbor may have good intentions in telling you about a study they found linking autism to the MMR vaccine. I can almost guarantee that your neighbor isn't aiming to harm children. They may not know that the paper they saw was written by the disgraced researcher Andrew Wakefield, nor that it was retracted and is now considered fraud by most of the medical community. Instead of assuming malice, consider that perhaps your neighbor made an honest mistake. Though in some cases you would be right, this principle will help you save face when you end up being wrong.

Admit fault

Humility is a virtue. We all make mistakes. When you do make a mistake, own it and admit fault swiftly and graciously. The first rule of arguing on the Internet seems to be “Always double down on your own stupidity,” but we strongly discourage this practice. It doesn’t move us forward, it wastes everyone’s time, and it jams up what could have been a productive venue for discussion. Such acts of defiance and childishness tarnish your credibility—and your credibility has to be worth more than the outcome of one argument.

Be clear

Imagine that you’re flying across the country and your seatmate engages you in a conversation about immigration or race relations or abortion or global warming. (We don’t recommend this, but it happens!) Your seatmate may be wrong; you know it; the person on the other side of them knows it; the flight attendant knows it. If there is any hope of convincing this individual, you need to be clear. A deluge of unordered facts never convinced anyone to abandon their previous beliefs. Your arguments need to be clear, understandable, persuasive, and preferably free of jargon. This usually requires as much effort, often far more effort, than it took to spot the bullshit in the first place.

For more serious rebuttals in your professional life, the presentation is as important as the argument. As we tell our students, effective refutation is hard work. For us, creating a clear presentation often requires creating and refining visualizations, developing a null model, creating a synthetic data set, testing the argument with friends, and double-checking to be sure we are not missing anything. The good news is that we can all get better at these things. The bad news is that we can’t possibly do this for every piece of bullshit that we find. Pick your battles. And when you have found your battle, make sure you win it. Do your homework in advance, and be clear in expressing your rationale.

Be pertinent

When we teach our students to call bullshit, we want to avoid creating a legion of “well, actually” guys. What’s a well-actually guy? It’s

the guy who interrupts a conversation to demonstrate his own cleverness by pointing out some irrelevant factoid that renders the speaker incorrect on a technicality.*

Let's take an example. I'm talking to a friend of mine over lunch, and I say: "It's interesting. There are lots of birds that trick other species into raising their offspring. Cuckoos, cowbirds, honeyguides, even some ducks. But mammals never do this. I wonder why not."

She thinks for a moment, and then observes, "I suspect it is because mammals don't lay eggs. That makes it a lot harder to sneak a baby in there!"

"Well, actually," I respond, "*some* mammals *do* lay eggs. Echidnas and platypuses, collectively known as monotremes and found only in Australia and Papua New Guinea, are both oviparous."

Let's face it, that was a totally annoying thing to do. And it had nothing to do with calling bullshit. Why not? Let's look at some of the factors that distinguish a caller of bullshit from a well-actually guy.

- **Relevance.** When you do a good job of calling bullshit, you undermine a claim in a fundamental way. Your objections negate the argument the speaker is trying to make. A well-actually guy doesn't move the discussion forward at all. Instead, he offers a pedantic or tangential objection that does not have much bearing on the core claims. When I objected about monotremes and eggs, I was not wrong, but my objection was almost entirely immaterial to the conversation we were having. There are a few egg-laying mammals, but my friend was right about 99.9 percent of mammal species.† In the end, my friend's statement "mammals don't lay eggs" isn't universally true. But her idea sounds right, and my objection does nothing to undermine it.

* "Hey, wait. 'Guy' is a gendered term," you object. "Actually, women sometimes do this too. . . ." They do. But in our experience, they do so far less often than men do. Congratulations. You're the "well, actually . . ." guy.

† Even the two exceptions conform to her rule. Though echidnas lay eggs, they never afford anyone an opportunity to slip in an egg of their own, because they lay them right in their own pouches. Platypuses don't have pouches, but the female seals herself into her tunnel before laying her eggs and doesn't emerge until after they have hatched. Again, there is no opportunity for anyone to slip in their own offspring.

- **Speaker's intention.** Typically, a caller of bullshit refutes someone who is intentionally bullshitting or being deceptive. A well-actually guy is different. He doesn't hesitate to contradict someone who is engaging in a dialogue in good faith. When I griped about monotremes, I fell on the wrong side of this line. My friend's suggestion was made in good faith. She wasn't trying to impress anyone. She wasn't trying to mislead anyone. She was simply trying to help me by answering my question.

- **Objector's motivations.** A well-actually guy doesn't care so much about where the argument is going as he does about demonstrating his own intellectual superiority. I didn't mention monotremes because they had any bearing on my friend's idea; I brought them up in order to make it clear that I'm better versed in matters zoological than she is—and possibly for the opportunity to use the word “oviparous” in conversation. Calling bullshit is not a way to signal your intelligence. Get a MENSA card if that sort of thing is important to you.

- **Audience.** When you call bullshit, you are usually trying to prevent a bullshitter or liar from misleading his audience. A well-actually guy doesn't care about protecting an audience, he is merely interested in demonstrating his own cleverness. When I objected about monotremes, no one else was even part of the conversation.

- **Power dynamics.** When you call bullshit, you may find yourself speaking truth to power. A well-actually guy isn't speaking truth to power. He's often punching down. His motivation is to put the speaker in her place while raising himself up.

- **Judiciousness.** A caller of bullshit makes a careful decision about whether it is worthwhile to speak up, derail a conversation, risk a confrontation, or make someone feel defensive. A well-actually guy simply cannot help himself. He hears something he believes he can contradict and doesn't have the self-control to think first about whether it is helpful to do so.

In the end, a well-actually guy has more in common with a bullshitter than he does with a caller of bullshit. A bullshitter disregards truth or logical coherence in order to impress or overwhelm an audience.

That's the well-actually guy. He doesn't care about advancing truth, or about the logical coherence of his objections. He is simply trying to impress or intimidate someone with his knowledge. Calling bullshit is not about making yourself look or feel smarter. If that is your goal, you are missing the point of this chapter, and indeed the whole point of this book. Effective bullshit calling is about making *others* smarter. This should be your barometer of success—and requires an extra level of social finesse.

Spotting bullshit is not easy, especially with the daily onslaught of misinformation. It takes practice and deliberate effort. Remember, calling bullshit is more than a claim. It's a powerful action, and can easily be misused. But if you make an effort to be clear and correct and to remain reasonably civil, most people will respect you.

Before you publicly refute a claim, ask yourself who your audience is—and whether that audience is worth your time. Some people are so entrenched in their beliefs that they will never be convinced, no matter how cogent the argument and how airtight the facts. Spend your time and effort on people who are willing to engage.

Above all, remember Neil Postman's dictum: "At any given time, the chief source of bullshit with which you have to contend is yourself." Confirmation bias can make us more confident than we ought to be, and humility is an important corrective. Self-reflection and an appreciation for the difficulty of getting to the truth: These are the marks of a mature thinker worth trusting. Sure, we want to keep the rest of the world honest—but for everyone's sake, let's start with ourselves.

Calling bullshit is more than a party trick, a confidence booster, or a way to sound impressive in front of your boss. It is a moral imperative. As we note in the opening line of the book, the world is awash with bullshit—from clickbait to deepfakes. Some of it is innocuous, some is a minor annoyance, and some is even funny. But a lot of the bullshit out there has serious consequences for human health and prosperity, the integrity of science, and democratic decision making.

The rise of misinformation and disinformation keeps us up at night. No law or fancy new AI is going to solve the problem. We all have to be a little more vigilant, a little more thoughtful, a little more careful when sharing information—and every once in a while, we need to call bullshit when we see it.

Acknowledgments

THIS BOOK COULD NOT HAVE BEEN WRITTEN WITHOUT THE HELP of many, many people.

First of all, we thank our wives, Holly Bergstrom and Heather West, for reading draft chapters and calling bullshit where necessary. Without their help, this book would have 83 percent more bullshit than it does. (We can get away with this one made-up statistic, because neither spouse will read this prior to publication.) We thank our children, Helen and Teddy, and Braylen and Camryn, for their patience when we were preoccupied with our writing, and for showing us that the ideas in this book do not require a college education to absorb.

We thank the hundreds of students and thousands of colleagues who have listened to us lecture about this material. Their attention, enthusiasm, questions, comments, suggestions, and challenges have been invaluable in helping us refine our message. We thank the Powers That Be at the University of Washington for initially tolerating and later enthusiastically promoting our vision for the Calling Bullshit course. Those we follow and those who follow us on Twitter have suggested numerous examples—many of which appear in this book—and they have kept us feeling just hip enough that we continue to embarrass ourselves in front of Generation Z. We are grateful to our friends who listened to us talk about bullshit *ad nauseum*, contributed their own views on the subject, and told us when we were pushing more bullshit than we were debunking. And we are grateful to our friends who helped us get away from the project at times, whether on the tennis courts or hiking trails.

Our friend, colleague, and coauthor, Jennifer Jacquet, provided the early encouragement we needed to even think of turning our ideas into a book, and provided valuable support along the way. Our agent, Max Brockman, helped us develop a loose set of ideas about bullshit into the concept for this book. Our editor at Random House, Hilary Redmon, took a poorly arranged draft and restructured it into the book you hold in your hands. Her editorial talents sliced out a hundred or more pages that you are blessed not to have read, and softened our natural dialect—tedious academic prose—into what you find here. Molly Turpin at Random House kept the entire project running and on schedule, while handling the logistical nightmares associated with producing over a hundred figures for a trade book. The beautiful text design you see here is the work of Barbara Bachman. Publicity manager London King, with assistance from marketer Ayelet Gruenspecht, guided us as we worked to reach a broad audience with the ideas we have developed in the book. Production manager Katie Zilberman, production editor Jennifer Rodriguez, and managing editor Rebecca Berlant somehow managed to pull the entire project together despite our best efforts to miss deadlines and make changes long after it was reasonable to even contemplate doing so. Our editor at Penguin Press, Casiana Ionita, shared our vision for the book from the start and was a wonderful source of enthusiastic support throughout the process. With her team—publicity manager Matt Hutchinson, marketer Julie Woon, production manager Sandra Fuller, and copy editor Stephanie Barrett—she developed the U.K. edition of the book. Pete Garceau designed the U.S. cover, and Richard Green designed the U.K. cover. Joel Klemenhagen, Matthew Brady, and the crew at The Shambles provided us with a home-away-from-home where we could discuss and develop the ideas in this book.

We are proud to serve as faculty members at one of the great public universities, the University of Washington. Our role at the University is not only to teach the students enrolled there. Our mission is to serve the people of our state, our country, and our world by helping to inform, educate, and cultivate the clarity of thought that will lead us all toward a world of greater truth and justice.