


-from: <http://www.salon.com/2017/06/18/explosive-growth-in-bullst-studies-the-latest-academic-frontier-in-the-age-of-you-know-who/>

“There is so much bullshit!” exclaims mathematical biologist Carl Bergstrom. “We’re drowning in it!” So begins the first session of **Calling Bullshit in the Age of Big Data**, a class Bergstrom is teaching with data scientist Jevin West at the University of Washington. It might be comforting if bullshit were confined to the realms of politics, PR and advertising. But that’s bullshit. It’s everywhere:

Science, my area, is conducted by press release more than it is carried out in the journals. Higher education — if you don’t know this by now, you guys are seniors — higher education rewards bullshit over analytic thought. Startup culture, we are here on the West Coast, startup culture has become sort of an elevation of bullshit to high art.

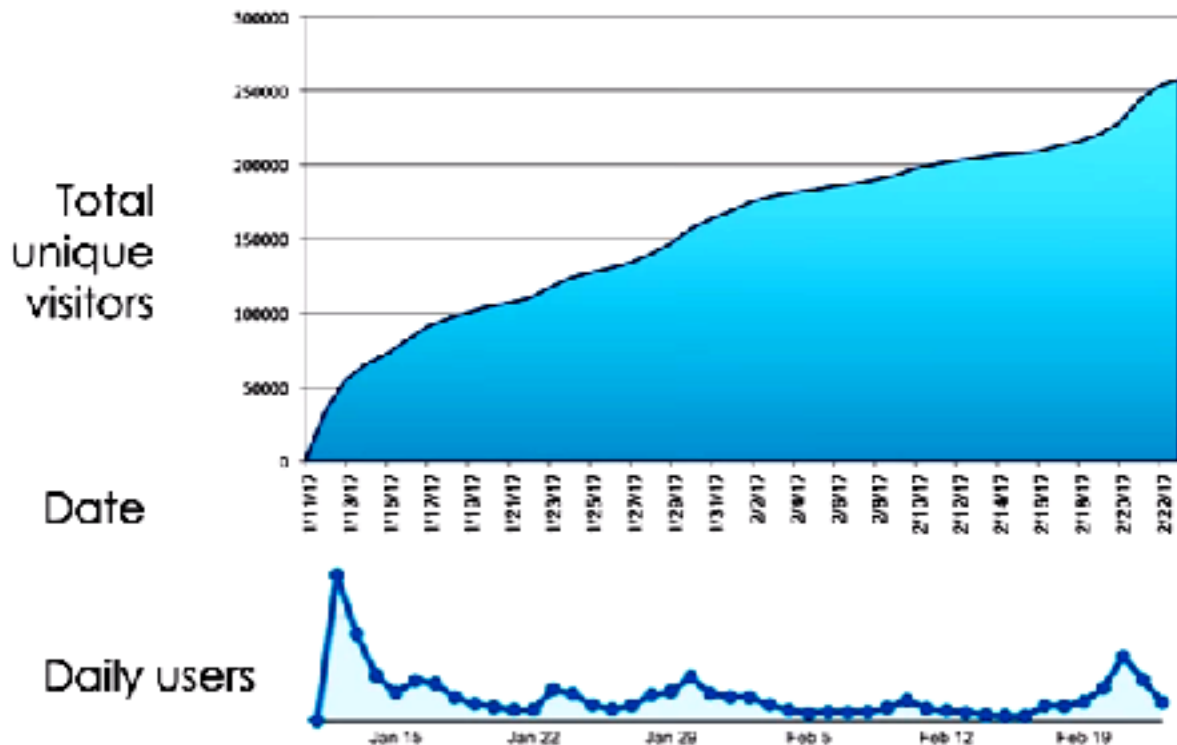
How much of an art? Well, **Calling Bullshit** does have its own website. As if it were a startup, when West took the stage he rolled out the following chart, reminiscent of an **infamous one** once used by Apple’s Tim Cook:

## Explosive growth at [callingbullshit.org](http://callingbullshit.org)



In response, Bergstrom called bullshit. “Sorry, I’m a little bit suspicious of what we just saw,” he said, going on to ask the class if anyone saw anything wrong. Someone quickly called out that it started at 100,000, not 0. But there’s a deeper problem, Bergstrom noted: Since it’s measuring *cumulative* unique visitors, it can only go up, never down. Is this growth really explosive? Viewed more realistically, in terms of daily visitors, not so much:

## Explosive growth at [callingbullshit.org](http://callingbullshit.org) ?



This playful exercise set an important tone for the course: The instructors are out to sharpen our senses to detect bullshit, wherever it may come from, not to stigmatize individuals, whose motives are often not knowable and could be innocent.

“I don’t want anyone to make anyone else feel unwelcome, let alone unsafe, for speaking their mind, or thinking differently than other people in the class,” Bergstrom said, talking about standards and expectations for the class. While the most malignant bullshitters may make everything personal, those who strive to free themselves from bullshit need to take the opposite approach. No one is immune, and bullshit can come from anywhere at all.

Indeed, much of Bergstrom and West's work is devoted to dealing with bullshit in science, which doesn't appear to come from anyone in particular, but from flaws in how science is done collectively — flaws a scientific study of science itself can help us understand and correct.

In this eclectic spirit, they present a wide array of clear, compelling illustrative examples. Class sessions — split into convenient bite-sized segments — have been **posted online**, with topics ranging from **basic bullshit-spotting** to **untangling correlation and causation** to **statistical pitfalls**, **big data deceptions**, **misleading data visuals**, and **problems plaguing science itself**.

The segment on **spurious correlations**, for example, presents a sampling of countless meaningless data-matches from **Tyler Vigen's website**, such as “Letters in winning word of Scripps National Spelling Bee” correlated with “Number of people killed by venomous spiders,” or “Number of people who drowned by falling into a pool” correlated with “Films Nicholas Cage appeared in.” But the same kind of spurious correlation crops up in “cutting-edge” research, trying to match genes with schizophrenia, for example.

The segment on **unfair comparisons** shows that a perennial internet favorite, ranking cities by crime rates, is overwhelmingly driven by how much or how little of a metro area is included in the central city.

The segment on what's called “**right censoring**” shows that a **widely circulated graph** of “Age of death and musical genre” tells us almost nothing about lifestyles and musician's mortality and everything about how old a cohort they belonged to.

The segment on **Brandolini's Bullshit Asymmetry Principle** helps explain why there's so much bullshit in the first place:

The amount of energy necessary to refute bullshit is an order of magnitude bigger than to produce it.

In the realm of data visualization, the segment on **misleading axes** presents a chart titled "Is truncating the Y-axis misleading?" It looks even, divided between "yes" and "no," but the chart itself is self-illustrating: It's truncated at 98 percent! The segment on **"glass slippers,"** meaning data shoved into forms that don't fit, includes nonsense examples like a "Periodic Table of Data [Science]," a "Microsoft Acquisitions and Investments" subway map and an "Internet Marketing Tree."

The big data segment "**Criminal Machine Learning**" deals with a recent paper, "**Automated Inference on Criminality Using Face Images,**" which had some people thinking of Steven Spielberg's film "Minority Report." It attempted to revive the long-discredited work of 19th-century physician Cesare Lombroso, who claimed to have identified physical criminal types, essentially as evolutionary throwbacks. As Bergstrom explained, "The idea of this paper is that Lombroso wasn't wrong, it's just that human eyes are too weak. If we throw fancy machine learning at it, we can rescue Lombroso and we can see the criminals." But the whole enterprise was biased by using arrestees' mugshots (whose subjects are likely to be frowning) as input, thus leading to a much simpler hypothesis: "This is actually a smile detector. It's not a criminality detector."

A segment on **reproducibility** described the infamous case of the erroneous claim that debt-to-GDP ratios over 90 percent spell doom for a country's economic growth. While eagerly embraced by austerity-minded elites, severely crippling economic recovery

from the Great Recession, the claim was demolished when a grad student got ahold of the data, and found serious data mistakes — including a simple Excel copying error. (I wrote about this [in 2013](#).) Though the whole world suffered from the error, at least it was uncovered. But that’s far too rare, West pointed out. “The problem is we don’t get rewarded for this in science.”

Amidst the wealth of silliness and seriousness just sampled, a few key insights really stand out, and the value of being able to own up to your own bullshit is one of the very first ones the course delivers.

## **Bullshit Defined**

But what is bullshit, anyway? And why does it merit a college class? There are a number of different plausible answers to the first question, including the following from the course’s [main web page](#):

Bullshit involves language, statistical figures, data graphics, and other forms of presentation intended to impress, overwhelm, or persuade — presented with a blatant disregard for truth, logical coherence, or what information is actually being conveyed.

Other definitions touched on are helpful as well. First, there’s Harry Frankfurt’s article (later book) “[On Bullshit](#),” commonly regarded as the starting point of modern “bullshit studies,” if there really is such a thing. Frankfurt distinguishes the bullshitter from the liar: the latter must care about the truth in order to hide it, whereas the former is utterly indifferent. As [I’ve quoted him before](#), Frankfurt wrote:

Bullshitters seek to convey a certain impression of themselves without being concerned about whether anything at all is true.

They quietly change the rules governing their end of the conversation so that claims about truth and falsity are irrelevant.

But Bergstrom and West draw attention to a significant refinement, “**Deeper into Bullshit**,” by G.A. Cohen, which shifts the focus away from the bullshitter and onto the bullshit instead. There’s also a subject-matter difference: Cohen sees Frankfurt focusing on everyday bullshit, while he’s more concerned with the academic kind. Forget the intention behind it, he says; what makes something bullshit when you encounter it? His short answer: “unclarifiable unclarity,” meaning discourse that “is not only obscure but which cannot be rendered unobscure,” and “where any apparent success in rendering it unobscured creates something that isn’t recognizable as a version of what is said.”

Read any good Foucault lately?

But Cohen also allows that there’s more to bullshit than that. He also refers to “rubbish, in the sense of arguments that are grossly deficient either in logic or in sensitivity to empirical evidence.” Which seemed a much better fit to the broad range of examples covered in the course, as well as reflecting the sense in the following quote, which provides an answer to the second question: Why does this merit a college class? The following, cited by West, comes from a speech by John Alexander Smith to Oxford students in 1914:

Nothing that you will learn in the course of your studies will be of the slightest possible use to you in the afterlife, save only this, that if you work hard and intelligently you should be able to detect when a man is talking rot, and that, in my view, is the main, if not the sole, purpose of education.

West cited that passage almost immediately after he stepped into his initial presentation. Call it what you will — rubbish, rot or

bullshit — developing the ability to detect it is the very essence of education, according to this view, which is in fact a venerable one. The essence of the Socratic method is, arguably, the exposure of bullshit. By all accounts, the ancient Greeks were drowning in it as well.

When I asked Bergstrom how he saw these different definitions fitting together, he first referenced their webpage definition, and then elaborated further:

According to this definition, the focus is really on impressing or overwhelming a listener/reader. I believe that this is a huge part of what statistics, machine learning algorithms, and complex data graphics do. They establish this veneer of authority, making themselves unquestionable through technical complexity and sophistication.

We like to talk about how statistical analyses or data science algorithms are like black boxes to the common reader: lots of data → black box (e.g., multinomial regression or random forest or whatever) → output (i.e., claims).

Without an advanced degree in stats or machine learning or some such, a reader is unable to “open the black box” (to borrow a metaphor from Bruno Latour) and thus the claim appears unassailable by such a reader. But we try to show our students that it’s not. If a claim is bullshit, it is rarely bullshit because something went wrong inside the black box (i.e., because of a technical problem with the analytic approach). Rather, it’s bullshit because the data going in are not what they claim to be (e.g., subject to selection biases, right censoring, cohort effects, or any number of other issues), or because the claims made are not actually supported by the immediate output of the black box.



In this way, we feel that almost all of our examples involve attempts to impress or overwhelm the audience with data and analysis.

## **Bullshit Detection**

A key aim of the course is to sharpen people's ability to make good ballpark judgments about what's happening *outside* the black boxes they encounter. One useful tool toward this end is what's known as **Fermi estimation**, after the famous Italian physicist Enrico Fermi, who had an uncanny knack for ballpark estimates — orders of magnitude. West used the example of a recent Fox News story, "Food stamp fraud at all-time high: Is it time to end the program?" arguing that \$70 million in reputed fraud meant that it was time to scrap the whole program.

"Seventy million dollars sounds like a lot of money," West said. But then he launched into Fermi estimations to get a handle on how large the program itself was. First, what fraction of Americans receive food stamps? Is it 1 percent, 10 percent or 100 percent? Who knows? But using Fermi estimation, almost everyone knows that 10 percent is around the right order of magnitude. Next, how much does the average recipient get? \$100? \$1,000? or \$10,000. Again, who knows the exact amount? But most folks correctly guess that \$1,000 is the right order of magnitude. And how many people in the United States? Using an order-of-magnitude scale, 300 million is the answer.

Put them all together, and the total spent is \$30 billion. And thus the fraud rate is *0.2 percent of expenditures*, which is pretty darned small. In fact, West noted, "If you ask anyone in retail, Nordstrom, Starbucks, any of the big companies ... they would

love the 0.2 percent. They're lucky if they're between 1.5 percent and 3 percent loss to fraud.”

That's just one example of how everyday people without any specialized training can develop double-checking habits to counter the flood of bullshit they encounter. But there's more to the course than that, Bergstrom pointed out:

Our course is about calling bullshit at least as much as it is about bullshit. As we note, calling bullshit has a broader scope of targets than bullshit alone: You can call bullshit on bullshit, but you can also call bullshit on lies, treachery, trickery or injustice.

Even if one subscribes to Frankfurt's or Cohen's definition of bullshit, virtually all of our examples are things that one can call bullshit upon.

## **Calling Bullshit In Science**

Arguably the most troubling — and most important — part of Bergstrom and West's course *Calling Bullshit* has to do with bullshit in science. First, in the segment “**P Values and the Prosecutor's Fallacy**,” Bergstrom explains why a common scientific metric — the p-value — can be so misleading. P-values are a measure of how likely a result would be, purely by chance. Published results usually require p-values of .05, or one in 20.

But that's not really what scientists ultimately want to know, as illustrated by the prosecutor's fallacy. A man is on trial for murder, and his DNA matches blood found at the scene. It could be someone else's, his lawyer argues, and the prosecutor scoffs: The chances of that are just one in a million, as if that means there's just a one-in-a-million chance he's got the wrong man. But that's looking at things the wrong way round. We *know* the DNA matches, and we know one in a million people match that blood,

so in the Seattle metro area — where Bergstrom teaches — there are 4 million people, *and thus four random matches*. So, the odds of the defendant's guilt (absent other evidence) are just one in five, as opposed to the odds of his innocence being one in 1 million. That's the prosecutor's fallacy.

The issue is taken up again in the session on problems in science, following a discussion of the growing replication problem, which has affected virtually every scientific field, as West describes. The problem is only occasionally related to individual failings, Bergstrom explains. Faking data is rare, and sloppiness not caught by reviewers is possible but not frequent. The vast majority of studies are “completely correct, completely right.” The problem is “**publication bias**,” which happens when authors and journals preferentially publish positive results. It's understandable, since negative results are boring. But the consequences are severe.

Ideally, scientists should be able to tell if a hypothesis is true or false by the cumulative distribution of studies. If the hypothesis is true, many more studies would reject the null hypothesis — say by 10 to 2. If it's false, only about 1 in 20 might support it. But because negative results are so rarely published either way, scientists often “can't distinguish which case we're in,” Bergstrom said. “That's a huge problem for trying to infer whether the null hypothesis is false by looking at multiple experiments,” People are developing statistical tools to try to compensate, but by “suppressing our negative results,” Bergstrom continued, “it's making it very hard for us to see whether our positives are these false positives that have occurred by chance or true positives that have occurred because something's really happening.”

This isn't a problem of individual bullshitting. It's not really a problem of bullshitting at all. But it most certainly *is* a problem of

bullshit. It's also a problem of scientifically understanding science. As Bergstrom explained in correspondence:

Science works well, but it could work even better. We're undermining our own progress with a set of norms [and] institutions that filter the information we are producing in ways that make it harder for us to reach correct conclusions about natural phenomena.

In the course, Bergstrom points out that science is done by humans, subject to our history and culture as well as biological nature. If science were done by bees, it would be very different. There's no way to escape nature, but we can become much more self-aware of it, as he suggested:

A fair bit of my recent research involves looking at the ways in which the historically contingent aspects of how science is organized (e.g., norms of credit assignment, concepts of what is publishable, structure of the academic career market, etc.) influence the things that science discovers, the things that it fails to discover, and the questions for which it gets the wrong answers. In other words, the social epistemology of science.

The openness to exploring how science itself can fail is a quintessential example of how science works differently from outside expectations. "Scientists make mistakes. Very bad!" we can just imagine *someone* tweeting. But it also needs to break from internal expectations too:

While science is in my opinion the greatest human invention, we would be mistaken to think of our scientific institutions and processes as the single optimal and inevitable methodology for coming to an understanding of the material universe. Doing so would blind us to the many inefficiencies in science and many opportunities for improvement.

The willingness to question science has costs as well, such as opening up new opportunities for bullshit. Bergstrom co-authored a paper, “**Publication Bias and the Canonization of False Facts**,” which was vulnerable on this account. “A prominent right-wing think tank released a commentary that misused our results to question climate science,” Bergstrom told me. “We were fortunate that this happened at the preprint stage,” which meant the authors could respond in the final published version.

The paper found that “unless a sufficient fraction of negative results are published, false claims frequently can become canonized as fact,” a further ramification of the prosecutor fallacy. But this *doesn't* throw all of science into question, as the authors explained:

Science denialists on both ends of the ideological spectrum might be tempted to invoke our findings as justification for their world-views. This would be a mistake. The facts that science denialists target are almost always very different from the types of facts we are modeling. We are modeling small-scale facts of modest import, the kind that would be established based on one or two dozen studies and then considered settled. The reality of anthropogenic climate change, the lack of connection between vaccination and autism, or the causative role of smoking in cancer are very different. Facts of this sort have enormous practical importance; they are supported by massive volumes of research; and they have been established despite well-funded groups with powerful incentives to expose any evidence that might give cause for skepticism. The process by which false claims can become canonized as fact in our model simply would not operate under these circumstances.

The length and carefulness of that passage reminds us of Brandolini's Bullshit Asymmetry Principle. Even *summarizing* a

refutation of bullshit can take more work and patience than bullshitting itself ever does.

## Why Do It?

As to why he and West created the course, taking away so much time from their other work, Bergstrom explained:

Our aim is not political in the sense of left- or right-wing, but it is civic. At the risk of sounding overly dramatic, I feel that people need to seek and be able to find reliable information in order for democracy to function.

He went on to quote an 1816 letter by Thomas Jefferson:

“If a nation expects to be ignorant and free in a state of civilization, it expects what never was and never will be. The people cannot be safe without information.” We simply want to help people of all political perspectives resist bullshit, because we are confident that together all of us can make better collective decisions if we know how to evaluate the information that comes our way.

There are intensified problems, he noted “in our so-called ‘post-truth society,’” pointing to a [tweet by Garry Kasparov](#): “The point of modern propaganda isn’t only to misinform or push an agenda. It is to exhaust your critical thinking, to annihilate truth.” He also cited Mark Galeotti’s [December 2016 op-ed](#) in the New York Times, arguing that “Americans should be taught the basic skills necessary to be savvy media consumers, from how to fact-check news articles to how pictures can lie.”

“We could not agree more,” Bergstrom concluded. “This is one thing Jevin and I actually know how to do, and once we realized that it was a near-obligation for us to devote a large fraction of our

time to this course and associated resources. This is why we have made the full course with all lecture videos, assignments and readings freely available to the world online.”