

## Not Even Scientists Can Easily Explain P-values



P-values have taken quite a beating lately. These widely used and commonly misapplied statistics have been blamed for giving a [vener of legitimacy to dodgy study results](#), encouraging [bad research practices](#) and promoting [false-positive study results](#).

But after writing about p-values again and again, and recently issuing a correction on a [nearly year-old story](#) over some erroneous information regarding a study's p-value (which I'd taken from the scientists themselves and [their report](#)), I've come to think that the most fundamental problem with p-values is that no one can really say what they are.

Last week, I attended the inaugural [METRICS conference at Stanford](#), which brought together some of the world's leading experts on meta-science, or the study of studies. I figured that if anyone could explain p-values in plain English, these folks could. I was wrong.

To be clear, everyone I spoke with at METRICS could tell me the technical definition of a

p-value — the probability of getting results *at least as* extreme as the ones you observed, given that the [null hypothesis](#) is correct — but almost no one could translate that into something easy to understand.

It's not their fault, said Steven Goodman, co-director of METRICS. Even after spending his "entire career" thinking about p-values, he said he could tell me the definition, "but I cannot tell you what it means, and almost nobody can." Scientists regularly get it wrong, and so do most textbooks, he said. When Goodman speaks to large audiences of scientists, he often presents correct and incorrect definitions of the p-value, and they "very confidently" raise their hand for the wrong answer. "Almost all of them think it gives some direct information about how likely they are to be wrong, and that's definitely not what a p-value does," Goodman said.

We want to know if results are right, but a p-value doesn't measure that. It can't tell you the magnitude of an effect, the strength of the evidence or the probability that the finding was the result of chance.

So what information can you glean from a p-value? The most straightforward explanation I found came from [Stuart Buck](#), vice president of research integrity at the Laura and John Arnold Foundation. Imagine, he said, that you have a coin that you suspect is weighted toward heads. (Your null hypothesis is then that the coin is fair.) You flip it 100 times and get more heads than tails. The p-value won't tell you whether the coin is fair, but it will tell you the probability that you'd get at least as many heads as you did if the coin was fair. That's it — nothing more. And that's about as simple as I can make it, which means I've probably oversimplified it and will soon receive exasperated messages from statisticians telling me so.

What I learned by asking all these very smart people to explain p-values is that I was on a fool's errand. Try to distill the p-value down to an intuitive concept and it loses all its nuances and complexity, said science journalist [Regina Nuzzo](#), a statistics professor at Gallaudet University. "Then people get it wrong, and this is why statisticians are upset and scientists are confused." You can get it right, or you can make it intuitive, but it's all but impossible to do both.

**Read more:** [Science Isn't Broken](#)

*Video by Jordan Schulkin.*