Smokers were never really protected from COVID, despite what early studies claimed

October 4 2021, by Mark Shrime

Credit: Unsplash/CC0 Public Domain

Early in the coronavirus pandemic, researchers stumbled on an unexpected finding: smokers seemed to be protected from COVID's worst effects. Initially discovered on a review of hospitalized patients in

China, this "smoker's paradox" was later reported in studies from [Italy](#) and [France](#).

But it turns out that this wasn't true, as a [massive study out of Britain](#) showed last month. Smokers were 80% more likely to be hospitalized than non-smokers. So what happened, and how did science get things so wrong?

The mathematician Pierre-Simon Laplace [once said](#): "The more extraordinary a fact is, the stronger proof it needs." The American cosmologist, Carl Sagan, famously reworded this as: "Extraordinary claims require extraordinary evidence." And, let's face it, for smokers, whose lungs get ravaged by tobacco, to have better outcomes in a respiratory disease is pretty miraculous.

Unfortunately, extraordinary proof is slow, complex and kind of boring. Public attention, on the other hand, is especially eager to latch on to the extraordinary.

Let's dissect what happened.

The first issue is that science is uncertain, a fact that makes us humans quite uncomfortable. Take a [weather forecast](#): if you're told there's a 10% chance of rain, you'll probably forgo the umbrella. I would. And nine out of ten times, I'd be right. But the other time, I'd regret my choices—and I'd complain about how wrong meteorologists can be.

The problem isn't meteorologists, though. It's my need for certainty. It's my subconscious translation of "there's a 10% chance of rain" into "it won't rain today."

This penchant is everywhere: in political polling, in presidential predictions—and even in doctors' visits. I want the doctor to tell me what

my sore throat is, not what it could be.

Everything is a probability

And that's how science works. Everything is a probability, and every new piece of information makes us update our probabilities. There's a famous example of this in statistics, first posed by the mathematician [Joseph Bertrand](#) (I promise I'll get back to the smoker's paradox in a second).

Say you have three identical boxes. One contains two gold coins, one contains two [silver coins](#), and the last contains one gold and one silver [coin](#). Pick one of the boxes at random (let's call it Box A). What are the chances that it has the two silver coins?

Exactly one-third.

Now, without looking in the box, take one coin out of it. If that coin is gold, what happens to the chance that Box A was the box that contained two silver coins?

It drops to zero. New information triggered an immediate probability update.

Which (finally) brings me back to COVID. In January 2020, we knew little about this virus. As good evidence trickles in, our probabilities update. It's why we're no longer sanitizing our mail but still recommending masks. No one can ever be 100% sure these recommendations are right—new evidence may emerge—but they reflect the best information we have.

The same goes with the smoker's paradox: before the pandemic, the evidence was that smoking did nothing good to your lungs. With

new—good—information, probabilities could have updated, shifting toward the extraordinary claim that smoking was protective.

And that's the second point: was this even good evidence?

It wasn't.

First, when they were reported, most papers on the smoker's paradox [had not been reviewed by other scientists (peer-reviewed)](). While a good number have gone on to peer-reviewed publication, [others have been retracted]() after it became clear that they had been funded by the tobacco industry. Pre-publication release is great for getting information out rapidly; it isn't great for making sure that information is sound.

Second, most of these studies [were small](). Although this isn't a death knell, it means that the evidence should be treated with caution. In other words, probabilities can update, just not a lot.

This makes intuitive sense: if you get 999 heads on 1,000 coin flips, you'd be pretty sure the coin was rigged. If you got two heads on three flips, you'd be a lot less sure. The studies suggesting the smoker's paradox had sample sizes in the teens to hundreds. The British study disproving it [had 421,000]().

Finally, and most subtly, the smoker's paradox studies asked a different question than they should have. They asked: "Of people currently in the hospital, how many smoke?" This is different from: "Compared with non-smokers, how likely are smokers in the population to be hospitalized?"

The first question looks at people who have already been admitted *and* have survived long enough to be studied. In other words, just like in Bertand's coin boxes, admission has already happened, and there are

many reasons that smokers weren't included in that group. Maybe they died faster than non-smokers, so weren't available to be counted. Maybe they were discharged to hospice at a different rate. The British study, on the other hand, studied the entire population, taking away this bias.

I'd argue, then, that science didn't get the smoker's paradox wrong. It was an interesting finding that led to a widely reported extraordinary claim. And if COVID teaches us nothing else, it should teach us to hold extraordinary claims—about smoking, vitamin D, zinc, bleach, gargling iodine, or nebulising hydrogen peroxide—to high standards.

Science moves slowly. Extraordinary claims do not. To paraphrase [Jonathan Swift](link), they fly along, while evidence comes limping after them.

This article is republished from [The Conversation](link) under a Creative Commons license. Read the [original article](link).

Provided by The Conversation