[Medical diagnosis fallacy]

Suppose you take a screening test for a medical condition that occurs in 1 in 10,000 people and, unfortunately, the screening test finds you to have the condition. You ask the doctor how accurate the screening test is – and unfortunately again, the doctor says that it's 99% accurate. Things seem pretty dire for you, but actually, things aren't as bad as they seem.

Let's reason through the probabilities at play here. We're concerned with the probability that you have the condition, given that you've tested positive to the screening test.

So there are actually two probabilistic events: Having the condition and testing positive on the test. Let's draw a tree diagram for these.

The probability that any given person has the condition, according to its frequency, is 1/10000, which means there's a 9999/10000 probability that any given person doesn't have the condition. Now from these events we look at what happens with the screening test. The screening test is 99% accurate, so for the people who do have the condition, 99/100 of them will be correctly told that they have tested positive, while 1/100 of them will be told that the screening test came up negative, even though they have the condition. For those 9999/10000 that don't have the condition, 99% of them will be told that they don't have the condition, which is true, but 1% of them will be told that they do have the condition, even though it's not true. So multiplying these probabilities through, we can work out what proportion of those that tested positive to the screening test actually have the condition. And we can see that because there are so many false positives, the likelihood of you actually having the condition after testing positive is less than 1%. This might seem counter-intuitive, but remember that there is only a 1 in 10000 chance of having the condition, so this probability represents a 100 fold increase.