T. Parent

Are you prone to the Base-Rate Fallacy?

You have been called to jury duty in a town where there are two taxi companies, Green Cabs Ltd. and Blue Taxi Inc. Blue Taxi uses cars painted blue; Green Cabs uses green cars. Green Cabs Ltd. dominates the market, with 85% of the taxis on the road.

On a misty winter night a taxi sideswiped another car and drove off. A witness says it was a blue cab. The witness is tested under conditions like those on the night of the accident, and 80% of the time she correctly reports the color of the cab that is seen. That is, regardless of whether she is shown a blue or a green cab in misty evening light, she gets the color right 80% of the time.

You conclude on the basis of this information:

(a) The probability that the sideswiper was blue is 80%
(b) It is more likely that the sideswiper was blue, but the probability is a little less than 80%.
(c) It is just as probable that the sideswiper was green as that it was blue.
(d) It is more likely than not that the sideswiper was green.

Pick an answer before reading on.

This question was devised by psychologists Amos Tversky and Daniel Kahneman. After extensive testing, they found that many people think that (a) or (b) is correct. Very few people think that (d) is correct. Yet (d) is the right answer!

People are prone to ignore that they are given *the probability that a randomly selected taxi will be green vs. blue.* This is the "prior probability," since it is known prior to the additional evidence (given by the witness). The prior probability here is also known as the *base rate*: It is the rate at which any given cab will be green vs. blue.

The fallacy is in ignoring the base rate when answering the question. The base rate effectively means that the witness says "blue" incorrectly more often than she says "green" incorrectly. And since she answered "blue," the numbers work out so that a green cab remains more likely.

In more detail: Assume for simplicity's sake that there are 100 cabs total, so that there are 85 green cabs and 15 blue cabs. Since the witness is 80% reliable, we'd expect the following:

1. She would correctly identify 80% of the green cabs, and 80% of the blue cabs. This would mean she would correctly identify 68 greens and 12 blues.
2. Conversely, it means she would misidentify 17 greens as "blue" and 3 blues as "green."

Thus, she would identify 29 cabs in total as "blue," but only 12 would be blue. So given that the witness says "blue," the chances of her being correct are 12/29, which is about 41%. It thus remains less than 50% likely that the cab was blue.

Here is an even more dramatic example to show why the base rate matters. (This example was mentioned in Forseman et al., ch. 6, but it's worth reviewing in more detail.)

There is a very rare and very terrible disease, "fictionitis." In the general population, only 1 person in 10,000 has the disease, but you decide to get tested anyway on a whim. The test is 99% accurate: If you have the disease, the test says YES with probability 99%—and if you do not have the disease, it says NO with probability 99%.

You get your results back and the test says YES. What is the probability that you have the disease?

(a) 99%
(b) Slightly less than 99%
(c) 50%
(d) About 1%

The answer in this case is also (d)—but the trick is to see *why* that is the right answer. The rough idea is this: The chances of having the disease is .01% for any member of the population, hence, .01% for any person that is tested. In contrast, the test gives the wrong answer to 1% of those people. So roughly, the likelihood of disease is still *100 times less* than the chances of a correct YES answer.

In more detail: If a total of 1,000,000 people are tested, then only 100 will have the disease. This represents the prior probability of having the disease. The reliability of the test then means that 99 of these 100 will get a YES result.

But what about the other 999,900 people who don't have the disease? The reliability of the test here means that 1% of 999,900 will get an incorrect YES result. But that ends up being *9999 people*!

So what are the chances that *your* YES result is correct? Well, there are 99+ 9999 = 10,098 people who got a YES result. But 9999 of 10,098 had an *incorrect* YES result. So the chances that your YES result is incorrect is 9999/10,098, which is approximately 99%!

Thus, even though the test is 99% reliable, it is still 99% likely that your YES result is incorrect. Since the base rate of the disease is *so* low, it ends up making it much more likely that your test result was a false positive, despite the high reliability of the test.