ELECTRONOTES
1016 Hanshaw Road
Ithaca, NY 14850

APPLICATION NOTE NO. 390

November 20, 2012

FUN WITH PROBABILITY PUZZLES 2 – FALSE POSITIVES

In the last AN we offered the "Monty Hall" puzzle as a popular example of a surprising result in probability theory. Here we want to take a look at a second case commonly brought forward and highly relevant – that of a false positive test. So let's consider some disease: the almost certainly fictional, but ever-threatening "Horrible Crud Disease" (HCD); a taunting tease of my childhood.

We suppose that there is a medical test for HCD. Like all medical tests, the test for HCD is not 100% accurate. In some cases, the test says you do have the disease when you don't (false positives) and on others, it says you don't have the disease when you do (false negatives). Nobody wants HCD anyway, so let's assume it is rare. So if you have just tested positive for HCD, what are the chances that you really do have it?

So make up some numbers. Let's assume that HCD is expected in 1% of the population. The test for HCD has a 5% false positive rate and a 2% false negative rate. So, you perhaps look at your positive HCD test and conclude that it could only be wrong, that is, a false positive (false alarm) in 5% of cases (one in 20). Isn't that what it means? Your consolation is that the HCD is known to be subject to 100% recovery in one day – so let's take out the anxiety element. Still, you are curious about the chances that you may not really have HCD.

So consider a population of 10,000 people. Of these, 1% or only 100 actually have HCD. The remaining 9900 people do not. Of the 100 people who do have HCD, 2 are missed because of the 2% false negative rate. Of the 9900 people who do not have HCD, 495 will nonetheless test positive due to the 5% false positive. Since there are 10,000 people total, 10,000 – 495 – 98 – 2 = 9405 don't have HCD and tested negative. Here is a table:

	Has HCD	Does Not Have HCD
Tests Positive	98	495
Tests Negative	2	9405

So let's look at just the people with positive tests. Of these 593 persons, only 98 or about 16.526% actually have HCD. So this completely changes the picture. Instead of feeling you were 95% sure to have HCD you now only have a chance of 16.5% of having it. Clearly the result seen here is the result of a fairly rare disease and a fairly high rate of false positives. Note that this table, well studied, tells us most everything we need.

Above we just constructed a table of possible outcomes and did not use any fancy notions of conditional probability and/or Bayes Formula. [More properly, arguably, we might have use an apostrophe and: Bayes' Formula, sometimes Bayes' Theorem.] While it is almost certainly true that a table of this sort is more valuable to the ordinary person than the "theorem", the notion of a mathematical theorem telling you that you are just fine might be very reassuring (or not).

Bayes Formula tells us how to modify an initial probability number in the light of all the evidence. Here is the formula:

$$P(D|T) = \frac{P(D)P(T|D)}{P(T)}$$

This powerful formula, while simple, is not that easy to understand without some study. It is just a multiplication of two things and then a division, but what are the symbols? First of all, they are all probabilities (P). The upright "bar" is a symbol that means "given". Thus P(D|T) means the probability that D is true given the fact that T is true (has occurred). [Caution: sometimes the opposite or even a different notation is used. In this case, D is the disease and T is a positive test.] This is exactly what we want to know: what is the probability that <u>you **have** the disease **given** that you **tested positive**.</u>

The point is that you might have mistaken the probability you want to know <u>for the one you were given</u>: P(T|D), the probability that <u>**if you have the disease**, the test will be **positive**.</u> We see that the two are skewed by the ratio P(D)/P(T). Here P(D) is the probability (before any tests) that you have the disease (the "background" rate). Accordingly any high probability of the test detecting the disease, P(T|D), which should be very high if the test is of any value at all (perhaps 90% to 99%) is reduced, first by a low value of P(D), the disease is not very prevalent, and further by a high value of false positives, which will make P(T) much larger, P(T) being all positive tests, true positives and false positives.

But, exactly what are the symbols in our example? P(D) is the background rate of the disease which we have taken to be 1%. This is an extremely important number to know, and it may be necessary to consider your exact current membership in a particular population. If you are a college student in classrooms of hundreds of students, influenza may be popular, but if you are overwintering at South Pole Station and have been there a couple of months with your fellow workers (isolated from visitors) even the "common cold" might be uncommon. Of course, you may be "immune" from certain gender-specific diseases if you are the wrong gender for that disease. So let's assume we do have some reasonable estimate of this "prior" probability P(D).

The probability P(T|D) is the reliability of the test in detecting the disease that is actually there. That is, we are concerned with false negatives. We have taken the rate of false negatives to be 2%. Thus P(T|D) is taken here to be 0.98. Again, we expect this number to be approaching one if the test is any good at all. That is, we want very few false negatives (misses), which may have serious consequences. This does not mean that too many false positives can be tolerated, and not just because they cause anxiety. In his 2008 book, *The Drunkard's Walk*, Leonard Mlodinow points out that <u>all</u> false negatives can be eliminated "simply" by gratuitously reporting all positive results!

Only the probability P(T) is likely to be difficult to obtain. This is the probability of a positive test, the true positives and the false positives. The true positives are P(D)P(T|D) and the false positives are [1 – P(D)] P(T|~D) = P(~D)P(T|~D). Here P(T|~D) is the rate of false positives, ~D standing for No Disease. We chose this to be 0.05. Thus P(T) = (0.01)(0.98) + (0.99)(0.05) = 0.0593. [Note: It's a sum of terms, and note that the major contribution to the 0.0593 is due to the 0.05 false positive rate.] The exact same number can be calculated directly from out table as (98 + 495)/10000.

Thus our result using Bayes Formula is:

$$P(D|T) = \frac{P(D)P(T|D)}{P(D)P(T|D) + \big(P(\sim D)\big)P(T|\sim D)}$$

$$= \frac{(0.01)(0.98)}{(0.01)(0.98) + (0.99)(0.05)}$$

$$= 0.16526$$

which is the same result we got playing with the tabulated numbers.*

These medical examples are common, and authors often point out the need to have medical professionals have a much better understanding of the often dramatic implications of conditional probability. Rightly so. Yet the mathematical considerations here extend to a much more general class of problems, and indeed to very general notions of Bayesian inference: revising our certainty, a "prior", as to whether a given hypothesis (H) is true or not in light of new evidence (E). The general formula is:

$$P(H|E) = \frac{P(H)P(E|H)}{P(E)}$$

We could well imagine using the formula in any case where we believe a result is expected with a certain associated probability, but, we have some additional information given that we have not used. At the very least, we need to think a bit about whether or not this additional information might affect our evaluation.

*   *   *   *   *   *   *

* This result was made up from scratch, but very similar examples are common in books and blogs. After doing my calculations I had in mind I needed to compare my results to two of these, one of which is a very real example which Mlodinow tells (with his characteristic good humor – surprising perhaps in this case – relating to his own false positive HIV test!). The other related to mammography in the book ***The Theory That Would Not Die*** (2011) by Sharon Bertsch McGrayne, an admirable general-reader discussion of Bayesian methods. There is very little math in this book – and what is there seems confused/wrong. The calculation relating to false positives is in Appendix B. The numbers there are wrong as the number 32/10,000 should be 32/40. It's the probability P(T|D), the true positives. (The final quoted result of 3% is correct.) Perhaps this typo is corrected in the paperback.