Notes from the Phil 3300 class - Thursday, Oct 16th, 2025

Bayes Theorem example:

What is the probability of getting three heads in three tosses of a fair coin?

Let's say HHH is the claim that you got three heads in a row.

The probability of getting three heads in a row is $P(HHH) = (1/2)^3 = 1/8$

What if you toss one, get a head, and then update? Lets say H1 stands for 'the first toss was heads'. Then P(HHH|H1) is the probability of getting three heads in a row given that the first toss was heads.

By Bayes's Theorem:

$P(HHH|H1) = P(H1|HHH) \times P(HHH) / P(H1) = 1 \times 1/8 / (1/2) = 1/4$

So the probability has changed from 1/8 to 1/4

If your first toss was a tail, then the new posterior would be 0.

$P(HHH|T1) = P(T1|HHH) \times P(HHH) / P(T1) = 0 \times 1/8 / (1/2) = 0$

Urn models:

One standard kind of probability problem uses urns (jars of marbles).

Imagine two urns - U1 is 80% red and U2 is 50% red. You draw a red ball from one of them. What are the odds it came from U1?

By Bayes's Theorem, $P(U1|R) = P(R|U1) \times P(U1)/P(R) = .8 \times P(U1) /P(R)$

So we need a prior probability on U1. If I had some prior reason to think it was U1 (or not), that would make a difference. Bets let's assume there are two urns and I chose randomly. Now $P(U1) = P(U2) = 1/2$

Now what about P(R)? This is the probability of drawing a red ball - but that is a *prior* probability. So not given any particular urn. In general, the probability of getting an outcome is a weighted average of the probability of getting that outcome in all possible scenarios weighted by the probability of those scenarios.

> **Comment:** If there are finitely many possibilities, then: $P(X) = \Sigma P(X|Bi) \times P(Bi)$. If infinitely many, it would be an integral and in general, this leads to an analytically

unsolvable problem. Which is related to why Bayesian statistics really starting going strong with the advent of computers.

So if there are two possibilities, U1 and U2, then P(R) = P(R|U1) x P(U1) + P(R|U2) x P(U2) = .8 x 1/2 + .5 x 1/2 = .65. In other words, since U1 and U2 are equally likely, the probability of drawing a red ball is the average of the two. It is either .8 or .5 so on average, it is .65. If you draw balls (with replacement) randomly over and over again, 65% of the time you would draw a red ball.

So going back to our original equation, P(U1|R) = P(R|U1) x P(U1)/P(R) = .8 x .5 /.65 = .4/.65 = 8/13 - approximately .615.

Intuitive answer:

We can figure out what the answer should be by thinking of natural frequencies. For example, imagine that there are 100 balls in each urn. Imagine the non-red balls are all blue (this doesn't matter for the problem). So U1 is 80R and 20B while U2 is 50R and 50B. So now of all the balls, 130/200 are red (that is where the P(R)=.65 comes from). If we draw a red ball, we know it is one of the 130 red balls. 80 of those balls come from U1. So the probability that it came from U1 is 80/130. This should make sense to you. But if we start to change the numbers - for example, if we are given that there are 80R 20B in U1 but 1R and 1B in U2, you have to manipulate the frequences to make them easy to combine again. Similarly, if P(U1) was 2/3 for example. Bayes' Theorem will get the correct answer in all cases without need for any creative thinking.

The taxicab problem:
This problem is answered in the extra reading from Hacking "Odd questions and Bayes' Rule".

If 85% of cabs in a town are green and the witness thinks she saw a blue cab and identifies the color of the cab correctly in 80% of cases regardless of color, then the posterior probability that the cab is blue is about ~41%. We can see that it should be less than 50% because the prior probability (15%) is lower than the false positive rate (20%). So in other words, it is more likely that the cab is green and she says blue [P(G&Wb) = P(G)xP(Wb|G) = .85 x .2] than it is that she the cab is blue and she got it right [P(B&Wb) = P(B) x P(Wb|B) = .15 x .8].

-- In class I also mentioned the odds ratio form of Bayes Theorem. This often leads to any easier way to see qualitative relationships (greater than or less than).

Imagine that we have two hypotheses H1 and H2.

In general, we know:

P(H1|O) = P(O|H1) x P(H1) / P(O) -- so P(O) = P(O|H1) x P(H1) / P(H1|O)

$P(H2|O) = P(O|H2) \times P(H2) / P(O)$ -- so $P(O) = P(O|H2) \times P(H2) / P(H2|O)$

By setting them equal to each other, $P(O|H1) \times P(H1) / P(H1|O) = P(O|H2) \times P(H2) / P(H2|O)$

By algebra:

$P(H1|O) / P(H2|O) = P(H1) / P(H2) \times P(O|H1) / P(O|H2)$

In other words, the ratio of the posteriors is the ratio of the priors multiplied by the ratio of the likelihoods. So the ratio of the likelihoods (called the Bayes factor) is a measure of the strength of the evidence - how much the probabilities are changing. For example, in the case of the urns, the odds (U1:U2) used to be 1:1. The ratio of the likelihoods is 8/5 so the new ratio of the posteriors is 8/5. For every 8 times you get a red ball from U1, there are 5 times you get a red ball from U2. A ratio of a:b is a probability of a/(a+b). So in this case, $P(U1|R) = 8/(8+5) = 8/13$.