Bayes' Theorem and Human Inference

This famous result is widely used as a mathematical formula to guide human inference.

We'll study the theorem in connection with a problem we have already looked at.

Failing a Drug Test

Your company institutes a drug testing program. Reliable evidence suggests that about 0.5% of the employees use cocaine. You have a very accurate test for cocaine. The sensitivity and specificity of the test are both .99. Your assistant Jason takes the test and gets a positive result. What is the probability that Jason uses cocaine?

Let C represent the event that Jason uses cocaine, and Pos represent the event that the cocaine test is positive.

We wish to know the conditional probability $\Pr(C \mid Pos)$.

Deriving Bayes' Theorem

We wish to know $\Pr(C \mid Pos)$. Earlier in the course, we learned the formula

$$\Pr(A \mid B) = \frac{\Pr(A \cap B)}{\Pr(B)}$$

From this we get

$$\Pr(C \mid Pos) = \frac{\Pr(C \cap Pos)}{\Pr(Pos)}$$

The problem is, *none of these quantities are given to us directly in the specification of the problem!!*

What We ARE Given

Let’s summarize what we are given in the statement of the problem:

$$\Pr(C) = .005$$

$$\Pr(\bar{C}) = 1 - \Pr(C) = .995$$

$$\Pr(Pos \mid C) = .99$$

$$\Pr(\overline{Pos} \mid C) = .01$$

$$\Pr(\overline{Pos} \mid \bar{C}) = .99$$

$$\Pr(Pos \mid \bar{C}) = .01$$

Deriving Bayes' Theorem

Here is where Bayes got very clever. Consider the numerator. We need $\Pr(C \cap Pos)$. We learned earlier in the course that there are two formulas we can use to compute that:

$$\begin{aligned}\Pr(C \cap Pos) &= \Pr(Pos)\Pr(C \mid Pos) \\ &= \Pr(C)\Pr(Pos \mid C)\end{aligned}$$

Notice that the *second* formula has two quantities we *are given* on the right-hand side. Bingo! Although we are not given $\Pr(C \cap Pos)$ directly, we can compute it.

But what about the denominator??

Deriving Bayes' Theorem

We need

$$\frac{\Pr(C \cap Pos)}{\Pr(Pos)}$$

and we now see how to compute the numerator. The denominator uses a really neat trick, using something we learned earlier in the course.

Deriving Bayes' Theorem

Recall that $Pos = \left(Pos \cap C\right) \cup \left(Pos \cap \bar{C}\right)$. That is, *all people who test positive are either Positive cocaine users or Positive non-users*! And they *cannot be both!* So, from the third axiom of probability theory,

So

$$\Pr(Pos) = \Pr(Pos \cap C) + \Pr(Pos \cap \bar{C})$$

But using the same trick we employed in the numerator, we get

$$\Pr(Pos) = \Pr(C)\Pr(Pos \mid C) + \Pr(\bar{C})\Pr(Pos \mid \bar{C})$$

All the quantities on the right-hand side are given!!

Deriving Bayes' Theorem

Putting everything together, we have

$$\Pr(C \mid Pos) = \frac{\Pr(C)\Pr(Pos \mid C)}{\Pr(C)\Pr(Pos \mid C) + \Pr(\bar{C})\Pr(Pos \mid \bar{C})}$$

$$= \frac{(.005)(.99)}{(.005)(.99) + (.995)(.01)}$$

$$= \frac{.00495}{.00495 + .00995} = \frac{99}{298}$$

The probability is slightly less than 1/3 that Jason is a cocaine user.

Deriving Bayes' Theorem

In general, then

$$\Pr(A \mid B) = \frac{\Pr(A)\Pr(B \mid A)}{\Pr(A)\Pr(B \mid A) + \Pr(\bar{A})\Pr(B \mid \bar{A})}$$

This formula is incredibly useful as a vehicle for *prescribing* how we should update our personal probabilities in the face of relevant evidence. However, there is another version of the formula that many people find much more useful for *understanding* the practical meaning of Bayes' theorem.

Odds vs. Probabilities

The Odds of A is the ratio of the probability that A occurs to the probability that it does not occur. That is,

$$\text{Odds}(A) = \frac{\Pr(A)}{\Pr(\bar{A})} = \frac{\Pr(A)}{1 - \Pr(A)}$$

We can move back and forth between odds and probabilities easily, because

$$\Pr(A) = \frac{\text{Odds}(A)}{1 + \text{Odds}(A)}$$

Odds-Likelihood Version of Bayes' Theorem

It is easy to prove that

$$\frac{\Pr(A \mid B)}{\Pr(\bar{A} \mid B)} = \frac{\Pr(A)}{\Pr(\bar{A})} \times \frac{\Pr(B \mid A)}{\Pr(B \mid \bar{A})}$$

or

Posterior Odds = Prior Odds $\times$ Likelihood Ratio

Example

The probability a woman in the general population has breast cancer is .003. The probability that a positive mammography result occurs given that a woman has cancer is .50. (This is called the sensitivity of the mammography test.) The probability that a negative result will occur given that a woman does not have breast cancer is .97. (This is called the specificity of the test.)

Suppose a woman has a mammography exam and gets a positive result. What is the probability that she actually has breast cancer?

The posterior odds are

$$\frac{.003}{.997} \times \frac{.50}{.03} = .05015$$, so the odds are still roughly 20 to 1 *against cancer!*

The posterior probability is only

$$\frac{.05015}{1+.05015} = .04776$$

If we use the other version of Bayes' theorem, we get

$$
\begin{aligned}
\Pr(C \mid Pos) &= \frac{\Pr(C)\Pr(Pos \mid C)}{\Pr(C)\Pr(Pos \mid C) + \Pr(\bar{C})\Pr(Pos \mid \bar{C})} \\
&= \frac{(.003)(.50)}{(.003)(.50) + (.997)(.03)} \\
&= \frac{.0015}{.0015 + .02991} = .04776
\end{aligned}
$$