Canonical example of Bayes' theorem in detail

John D. Cook*

February 8, 2008

The most common elementary illustration of Bayes' theorem is medical testing for a rare disease. The example is almost a cliché in probability and statistics books. And yet in my opinion, it's usually presented too quickly and too abstractly. Here I'm going to risk erring on the side of going too slowly and being too concrete. I'll work out an example with numbers and no equations before presenting Bayes theorem. Then I'll include a few graphs.

Suppose there are 1,000,000 people in a given population, and 1,000 of these people carry a certain genetic mutation. And suppose there's a test for this mutation that is 95% accurate in the sense that 95% of those who have the mutation will test positive, and only 5% of those who do not have the mutation will test positive. In other words, the false positive rate and the false negative rate are both 0.05. (In general these rates will differ, but they're the same in this example for simplicity.)

You take a test for the mutation and the result is positive. What is the probability that you really have the mutation? Your first response may be 95%, but that's the probability that the test will be positive given that you have the mutation. You want to know the opposite, the probability of having the mutation given that the test is positive. In symbols, we know

$$\Pr(\text{positive} \,|\, \text{mutation}) = 0.95$$

but we want to know

$$\Pr(\text{mutation} \,|\, \text{positive}).$$

Here's a quick example to show that $\Pr(A \,|\, B)$ and $\Pr(B \,|\, A)$ may be very different. Suppose you pick an American citizen at random and want to know the probability that this person is a former U.S. president. To date, all U.S. presidents have been male, but the vast majority of males have not been president.

*http://www.JohnDCook.com

So

$$\Pr(\text{male} \mid \text{former president}) = 1$$

but

$$\Pr(\text{former president} \mid \text{male}) \approx 0.$$

Back to our mutation test. Suppose all 1,000,000 people were tested for the mutation. How many people would test positive? Ninety-five percent of the 1,000 who have the mutation and five percent of the 999,000 who do not, for a total of $0.95 \times 1{,}000 + 0.05 \times 999{,}000 = 50{,}900$. How many of those who tested positive had the mutation? 950. If you test positive, it is far more likely that you're one of the 49,950 who tested positive but did not have the mutation than that you were one of the 950 who had the mutation and tested positive. In other words, a person with a positive results is more likely to be among the 49,950 false positives than the 950 true positives.

The probability that you have the mutation given that you tested positive is $950/(950 + 49{,}950)$ or about 2%. So in this example

$$\Pr(\text{positive} \mid \text{mutation}) = 0.95$$

but

$$\Pr(\text{mutation} \mid \text{positive}) = 0.0187$$

Knowing that you tested positive increased your probability of having the mutation from 0.001 to 0.0187, but not all the way to 0.95.

Let's go back and look at what would happen if 40% of the original population had the mutation. If everyone were tested, 95% of the 400,000 who have the mutation would test positive, and 5% of the 600,000 who do not have the mutation would test positive. Your probability of having the mutation given that you tested positive would be 380,000 / (380,000 + 30,000) or about 93%. Since the mutation is common, the number of false positives (30,000) is fairly small compared to the number of true positives (380,000), and so a positive is likely to be a true positive.

Notice that in both examples, a positive test result increases your probability of having the mutation, no matter how common the mutation is. But the size of the increase depends greatly on the prevalence of the mutation. When the mutation was rare (0.1%), the *relative* increase in the probability of having the mutation due to a positive test result was large ($0.0187/0.001 = 18.7$), but the *absolute* increase in the probability was small ($0.0187 - 0.001 = 0.0177$). When the mutation rate was large (40%), the relative increase in the probability of

mutation was fairly small ($0.93/0.40 = 2.325$) but the absolute increase was large ($0.93 - 0.40 = 0.53$).

Now we break out Bayes' theorem. In it's simplest form, Bayes theorem says

$$\Pr(H \mid E) = \frac{\Pr(E \mid H)\Pr(H)}{\Pr(E)} = \frac{\Pr(E \mid H)\Pr(H)}{\Pr(E \mid H)\Pr(H) + \Pr(E \mid \neg H)\Pr(\neg H)}.$$

Bayes' theorem is true for any events E and H but E suggests "evidence" and H suggests "hypothesis." (Bayes' theorem is much easier to prove than to understand. To prove the theorem, just apply the definition of conditional probability and simplify.) In our application, we start with $\Pr(E \mid H)$, the probability of the evidence (the test result) given the hypothesis (presence of the mutation), and we use Bayes' theorem to compute $\Pr(H \mid E)$, the probability of the mutation given the test result.

In our application, the numerator in Bayes' theorem is the sensitivity of the test (ratio of positive test results to total number of people with the mutation) multiplied by the prevalence (the proportion of the total population with the mutation). The denominator is the probability of a positive test result.

When the prevalence of the mutation is 0.001, Bayes' theorem tells us the probability of having the mutation given a positive test result is

$$\frac{0.95 \times 0.001}{0.95 \times 0.001 + 0.05 \times 0.999} = 0.0187.$$

When the prevalence of the mutation is 0.40, the probability of having the mutation given a positive test result is

$$\frac{0.95 \times 0.4}{0.95 \times 0.4 + 0.05 \times 0.6} = 0.93.$$

Next we include a graph to show how the probability of mutation given a positive test varies as a function of prevalence. The horizontal line at 0.95, the probability of a positive test given the mutation, is added for reference. Note that for prevalence less than 0.526,

$$\Pr(\text{mutation} \mid \text{positive}) < \Pr(\text{positive} \mid \text{mutation})$$

but for larger values of prevalence the opposite is true.

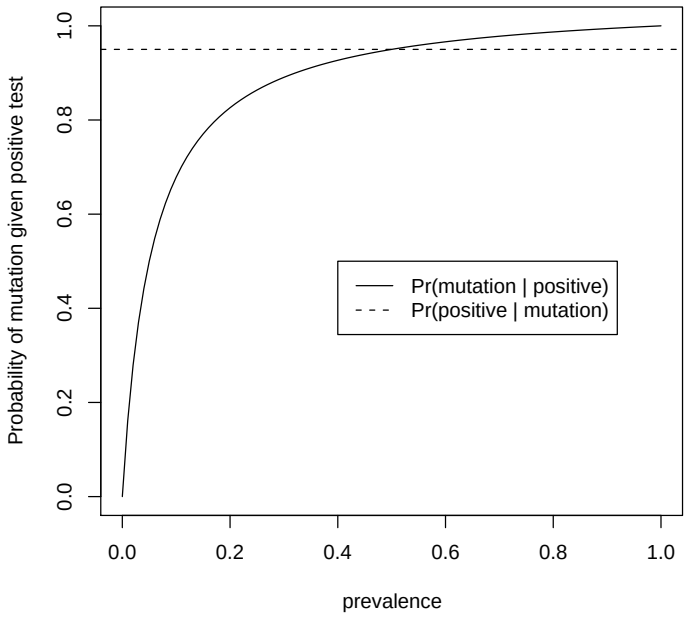

Now let's look more closely at the relative and absolute changes in the probability of having the mutation as the prevalence changes. Let p be the prevalence of the mutation, the unconditional probability that someone has the mutation. Bayes theorem says that the probability of having the mutation given a positive test result is $95p/(90p+5)$. So the relative increase in the probability of having the mutation due to a positive test result is $R(p) = 95/(90p+5)$. The absolute change is $A(p) = 95p/(90p+5) - p$. The following graphs plot $R(p)$ and $A(p)$.

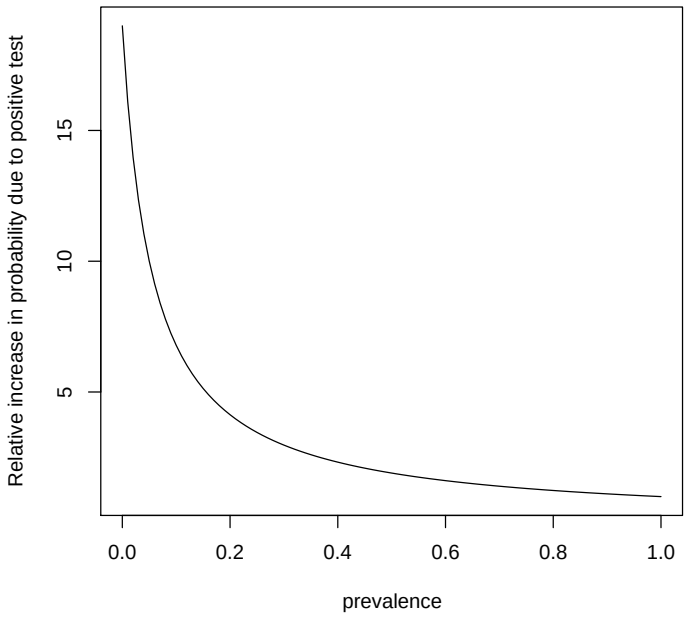
Relative increase in probability due to positive test
5
10
15
0.0
0.2
0.4
0.6
0.8
1.0
prevalence

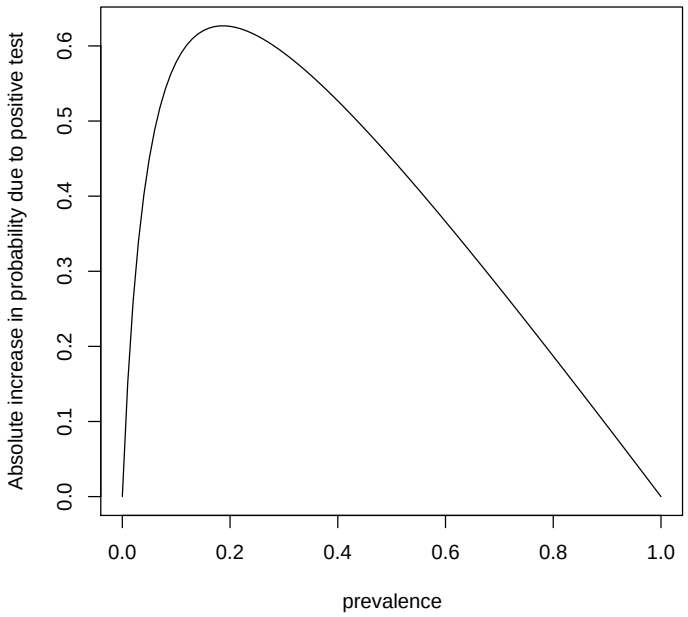

As our particular examples with $p = 0.001$ and $p = 0.40$ suggest, the relative increase in probability of mutation due to a positive test result decreases as prevalence increases. The absolute increase is small when the prevalence is small (or very large) but increases for moderate levels of prevalence, reaching a maximum around $p = 0.187$.