Some notes on Bayes' Theorem

Murray Gerstenhaber

Prepared for the exclusive use of students in the course *Statistical Methods for Lawyers* at University of Pennsylvania Law School. Copyright 2004 by Murray Gerstenhaber. PRELIMINARY DRAFT

There is a report of a death in a remote area. The reported symptoms of the victim indicate that cause could have been Lhasa fever or pneumonia. What are the chances that it was the latter?

A test to detect a certain disease sometimes fails to detect it and sometimes produces a false positive report. The test on a certain patient was positive. What are the chances that he really has the disease?

You are blindfolded before a table on which there are two urns, each containing a mixture of black and white balls but with different proportions. You pick a ball and discover that it is black. What are the chances that it came from the urn on the right?

These are all the same problem: the first two examples are just specific cases of the third. In the first the two urns are "Lhasa fever" and "pneumonia", each containing a white ball for every area victim of the particular disease who survived and a black one for every one that died. In the second the urns are "tested and really has the disease" and "tested but doesn't have the disease"; the black balls in each urn are positive test results.

Obviously we can't answer the question without more information, so let's suppose, in the urn problem, that these are the data: The urn on the left contains 100 balls of which 75 are black, while the right one contains only 50 balls of which 10 are black. Suppose also that you are right-handed and that your tendency would be to pick the urn on the right two-thirds of the time. What can we say now? In the disease problem, this would correspond to pneumonia being twice as prevalent as Lhasa fever in the area where the death was reported, with a mortality rate of 75% for Lhasa fever and only 20% for pneumonia.

Returning to the urns, suppose that each ball in addition to being black or white has a mark indicating from which urn it came and that we now combine all the balls in one large urn (and mix well). Forgetting about the colors, if we were to draw a ball at random from this larger urn (which now contains twice as many balls from the left urn as from the right)the chances are that two times out of three it would come from the left urn. This does not reflect our assumptions. In probability terms, the probability that the ball chosen would come from the left urn would be 2/3 and that it would come from the right would be 1/3, while we actually want the reverse. So we can't simply mix all the balls. One way to get the correct probabilities is to add balls to the right urn keeping the ratio of black to white balls the same. Originally one in five of the balls was black. Putting 200 balls in the right urn of which 40 are black wouldn't change the chances of drawing a black ball if we picked one at random from that urn. Now if we mix the contents of the two urns in a larger one and draw a ball at random it will indeed happen that 2/3 of the time the ball we choose will have come from the urn on the right, which is just what we wanted.

At this point it will make no difference if we have two urns on the table or just the larger combined one (in which the balls, in addition to being black or white, are still marked "left" or "right", so we can still tell where they came from). Now imagine picking a ball at random from the large urn and finding that it was black. What is the probability of that it came from the right urn? The white balls no longer matter (we didn't pick one of them), only the black balls count. The total number of black balls in the large urn was 115, of which 75 were from what was the left urn and 40 from the right. Therefore, the probability that the black ball we picked was from the right urn is 40/115. In the disease model this would say that even though the prevalence of pneumonia in the reporting area was twice as great as that of Lhasa fever, the greater mortality rate for Lhasa fever brings the probability that it was the cause of the reported death up to 75/115.

This example illustrates some basic ideas. Originally we had certain probabilities of picking from each of the two urns, 1/3 for the left urn and 2/3 for the right. These are the *a priori* probabilities. Then some new information emerged: the ball actually chosen was black. This can not tell us with certainty into which urn we put our hand (unless one of the urns has no black balls at all) but it does change the probability of our choice. The probability that we actually chose from the left urn is now 75/115 and from the right is 40/115. These are the *a posteriori* probabilities. Of course it was not really necessary to mix the balls from the urns. Having adjusted the numbers so that the totals in the two urns were proportional to the *a priori* probabilities, the numbers of black balls in the urns have become proportional to the *a posteriori* probabilities (assuming that when our blindfold was removed we saw that we had chosen a black ball).

We will develop a simple formula for how the probabilities have changed, but for the moment observe that the simple idea we used can carry us much further. First, we aren't limited to two urns. Suppose we were blindfolded in front of five urns, each containing a mixture of black and white balls, and we knew the *a priori* probability of picking from each of the urns. (There might be five diseases whose symptoms fit those reported for the mysterious death.) Adjust the number of balls in each urn (keeping the proportions of black and white in each fixed)so that the total numbers in the urns are proportional to their *a priori* probabilities. If it turns out that you picked a black ball then the numbers of black balls in the urns will be proportional to the *a posteriori* probabilities. Second, we didn't have to limit ourselves to black and white. There could also be some red balls in the urns. (In the disease model white might correspond to total recovery, black to death, and red to survival but with some morbidity or permanent injury, for example, a lingering paralysis in a case of polio). We could have numerous urns and numerous colors. Again, once the numbers of balls in the urns are adjusted to reflect the *a priori* probabilities of choosing from each of the urns (keeping the proportions of balls of each color fixed within each urn) then for each color the numbers of balls of that color in the urn will be proportional to the *a posteriori* probabilities.

Finally, let's return, for simplicity, to the case of the two urns with just black and white balls. Suppose that we draw a ball, that the referee sees that it is black, that he returns it to urn from which it came (and stirs), and that while keeping our blindfold on and *our hand in the same urn* we draw a second ball and see that it is again black. (In the disease case, suppose someone in the same household as the first victim has been stricken, presumably through direct transmission, and has also died.) What are the *a posteriori* probabilities? Here all we need to do is to take the *a posteriori* probabilities for the first drawing as the prior ones for second.

The method of adjusting the numbers of balls in the urns to reflect the *a priori* probabilities can rapidly produce unwieldy numbers. At the expense of introducing fictional fractional balls one can develop a simple formula for the *a posteriori* probabilities. Suppose, to fix the ideas, that there are five urns, $U_1, U_2, \ldots, U_5$, and that standing blindfolded in front of them you would initially have had a probability which we will write as $p(U_1)$ of picking the ball from the first, $p(U_2)$ for the second, and so forth. (Of course, 5 could be replaced by any number. In the initial example there were two urns U_1 and U_2 with probabilities $p(U_1) = 1/3, p(U_2) = 2/3$.) These are the *a priori* probabilities. They are non-negative numbers which must add up to 1 since these are all the possibilities that are allowed and no two can occur simultaneously; you will draw a ball from one and only one of the five urns. Suppose now that when the urns are mixed together that the first urn contributes the fractional number $p(U_1)$ balls to the large urn, the second contributes the fractional number $p(U_2)$, and so forth. It all adds up to one ball in the large urn, but that one can be divided as finely as we like. (If this is hard to visualize, suppose that the probabilities $p(U_1), \ldots, p(U_5)$ are expressed as percentages – whole numbers, no fractions. These must add to 100%, so that all told we have 100 balls in the five urns. In the initial example we would have approximately 33 balls in the left urn and 67 in the right one. If you want more accuracy think of 333 balls in the left urn and 667 in the right. It is clearly better to abandon this and just think of $p(U_1)$ balls in the first urn, and so forth.) Now *supposing that you have put your hand into the first urn* we denote the probability the probability that the ball you draw is black by $p(B|U_1)$, read "*the probability of* B (a black ball) *given* (that your hand is in urn) U_1. In the initial example with 75 out of the hundred balls in the first urn being black, this is .75. We may do the same for the other urns, however many there are. In the initial example, $p(B|U_2) = 10/50 = .2$. Since the first urn contributes $p(U_1)$ balls to the total mix, and of these the fraction that are black is $p(B|U_1)$, the number of black balls the first urn contributes to the total mix is $p(U_1)p(B|U_1)$. In the initial example this would be $1/3 \times$

.75 or .25 black balls. Similarly, the second urn would contribute $2/3 \times .2$ or .133 black balls. The mix would therefore contain $.25 + .1333 = .3833$ black balls. Knowing that a black ball has been picked, the probability that it came from the first urn (in this example) is therefore .25/.3833 or approximately .65. (This is very close to 75/115; the error is due to rounding.) In the case where we have five urns, the total number of black balls they will contribute to the mix will be $p(U_1)p(B|U_1) + \cdots + p(U_5)p(B|U_5)$. Of these the number that will have come from the first urn is $p(U_1)p(B|U_1)$. Suppose you find that you have drawn a black ball. The *a posteriori* probability that it came from the first urn is denoted $p(U_1|B)$, read the *probability of* (having picked from urn) U_1 *given* (that the ball is) B(lack). We can now compute this easily; the result is **Bayes' Theorem**:

$$p(U_1|B) = \frac{p(U_1)p(B|U_1)}{p(U_1)p(B|U_1) + \cdots + p(U_5)p(B|U_5)}.$$

A similar formula gives the *a posteriori* probability that the black ball came from the second urn (replace U_1 by U_2) and likewise for any of the urns. Also, "5" can be replaced by any number (at least 2). Notice that the *a posteriori* probabilities also sum to 1 (as they must). Also, it really did not matter how many colors there were; "B" might stand for any of them.

In the special case where there are only two urns one can think of this formula in the somewhat simpler way that is given in Chapter 1 of the notes for this course.