From: I. Hacking, An Introduction to Probability and Inductive Logic (Cambridge: 2001)

7 Bayes' Rule

One of the most useful consequences of the basic rules helps us understand how to make use of new evidence. Bayes' Rule is one key to "learning from experience."

Chapter 5 ended with several examples of the same form: urns, shock absorbers, weightlifters. The numbers were changed a bit, but the problems in each case were identical.

For example, on page 51 there were two urns A and B, each containing a known proportion of red and green balls. An urn was picked at random. So we knew:

Pr(A) and Pr(B).

Then there was another event R, such as drawing a red ball from an urn. The probability of getting red from urn A was 0.8. The probability of getting red from urn B was 0.4. So we knew:

Pr(R/A) and Pr(R/B).

Then we asked, what is the probability that the urn drawn was A, *conditional* on drawing a red ball? We asked for:

Pr(A/R) = ? Pr(B/R) = ?

Chapter 5 solved these problems directly from the definition of conditional probability. There is an easy rule for solving problems like that. It is called *Bayes' Rule.*

In the urn problem we ask which of two *hypotheses* is true: Urn A is selected, or Urn B is selected. In general we will represent hypotheses by the letter H.

We perform an *experiment* or get some *evidence*: we draw at random and observe a red ball. In general we represent evidence by the letter E.

Let's start with the simplest case, where there are only two hypotheses, H and ~H. By definition these are mutually exclusive, and exhaustive.

Let E be a proposition such that Pr(E) > 0. Then:

$$Pr(H/E) = \frac{Pr(H)Pr(E/H)}{Pr(H)Pr(E/H) + Pr(\sim H)Pr(E/\sim H)}$$

This is called *Bayes' Rule* for the case of two hypotheses.

PROOF OF BAYES' RULE

$$Pr(H\&E) = Pr(E\&H)$$

$$\frac{Pr(H\&E)Pr(E)}{Pr(E)} = \frac{Pr(E\&H)Pr(H)}{Pr(H)}$$

Using the definition of conditional probability,

$$Pr(H/E)Pr(E) = Pr(E/H)Pr(H).$$

$$Pr(H/E) = \frac{Pr(H)Pr(E/H)}{Pr(E)}$$

Since H and (~H) are mutually exclusive and exhaustive, then, by the rule of total probability on page 59,

$$Pr(E) = Pr(H)Pr(E/H) + Pr(\sim H)Pr(E/\sim H).$$

Which gives us Bayes' Rule:

$$(1) \quad Pr(H/E) = \frac{Pr(H)Pr(E/H)}{Pr(H)Pr(E/H) + Pr(\sim H)Pr(E/\sim H)}$$

GENERALIZATION

The same formula holds for any number of *mutually exclusive* and *jointly exhaustive* hypotheses:

$H_1, H_2, H_3, H_4, \ldots, H_k$, such that for each i, $Pr(H_i) > 0$.

Mutually exclusive means that only one of the hypotheses can be true. *Jointly exhaustive* means that at least one must be true.

By extending the above argument, if Pr(E) > 0, and for every i, $Pr(H_i) > 0$, we get for any hypothesis H_k,

$$(2) \quad Pr(H_j/E) = \frac{Pr(H_j)\ Pr(E/H_j)}{\Sigma[Pr(H_i)\ Pr(E/H_i)]}$$

Here the Σ (the Greek capital letter sigma, or S in Greek) stands for the *sum* of the terms with subscript i. Add all the terms $[Pr(H_i)Pr(E/H_i)]$ for i = 1, i = 2, up to i = k.

Formula (1) and its generalization (2) are known as Bayes' Rule.

The rule is just a way to combine a couple of basic rules, namely conditional and total probability. Bayes' Rule is trivial, but it is very tidy. It has a major role in some theories about inductive logic, explained in Chapters 13–15 and 21.

URNS

Here is the urn problem from page 51:

Imagine two urns, each containing red and green balls. Urn A has 80% red balls, 20% green, and Urn B has 60% green, 40% red. You pick an urn at random, and then can draw balls from the urn in order to guess which urn it is. After each draw, the ball drawn *is replaced*. Hence for any draw, the probability of getting red from urn A is 0.8, and from urn B it is 0.4.

$Pr(R/A) = 0.8 \qquad Pr(R/B) = 0.4 \qquad Pr(A) = Pr(B) = 0.5$

You draw a red ball. What is P(A/R)?

Solution by Bayes' Rule:

$$Pr(A/R) = \frac{Pr(A)Pr(R/A)}{Pr(A)Pr(R/A) + Pr(B)Pr(R/B)}$$

$$= (0.5 \times 0.8)/[(0.5 \times 0.8) + (0.5 \times 0.4)] = 2/3.$$

This is the same answer as was obtained on page 51.

SPIDERS

A tarantula is a large, fierce-looking, and somewhat poisonous tropical spider.

Once upon a time, 3% of consignments of bananas from Honduras were found to have tarantulas on them, and 6% of the consignments from Guatemala had tarantulas.

40% of the consignments came from Honduras. 60% came from Guatemala.

A tarantula was found on a randomly selected lot of bananas. What is the probability that this lot came from Guatemala?

Solution

Let G = The lot came from Guatemala. Pr(G) = 0.6.
Let H = The lot came from Honduras. Pr(H) = 0.4.
Let T = The lot had a tarantula on it. Pr(T/G) = 0.06. Pr(T/H) = 0.03.

$$\Pr(G/T) = \frac{\Pr(G)\Pr(T/G)}{\Pr(G)\Pr(T/G) + \Pr(H)\Pr(T/H)}$$

Answer: Pr(G/T) = (.6 × .06) / [(.6 × .06) + (.4 × .03)] = 3/4

TAXICABS: ODD QUESTION 5

Here is Odd Question 5.

> You have been called to jury duty in a town where there are two taxi companies, Green Cabs Ltd. and Blue Taxi Inc. Blue Taxi uses cars painted blue; Green Cabs uses green cars.
>
> Green Cabs dominates the market, with 85% of the taxis on the road.
>
> On a misty winter night a taxi sideswiped another car and drove off. A witness says it was a blue cab.
>
> The witness is tested under conditions like those on the night of the accident, and 80% of the time she correctly reports the color of the cab that is seen. That is, regardless of whether she is shown a blue or a green cab in misty evening light, she gets the color right 80% of the time.
>
> You conclude, on the basis of this information:
>
> ______(a) The probability that the sideswiper was blue is 0.8.
> ______(b) It is more likely that the sideswiper was blue, but the probability is less than 0.8.
> ______(c) It is just as probable that the sideswiper was green as that it was blue.
> ______(d) It is more likely than not that the sideswiper was green.

This question, like Odd Question 2, was invented by Amos Tversky and Daniel Kahneman. They have done very extensive psychological testing on this question, and found that many people think that (a) or (b) is correct. Very few think that (d) is correct. Yet (d) is, in the natural probability model, the right answer! Here is how Bayes' Rule answers the question.

Solution

Let G = A taxi selected at random is green. Pr(G) = 0.85.
Let B = A taxi selected at random is blue. Pr(B) = 0.15.
Let W_b = The witness states that the taxi is blue.
$\Pr(W_b/B) = 0.8$.

Moreover, $\Pr(W_b/G) = 0.2$, because the witness gives a *wrong* answer 20% of the time, so the probability that she says "blue" when the cab was green is 20%.

We require $\Pr(B/W_b)$ and $\Pr(G/W_b)$.

$$\Pr(B/W_b) = \frac{\Pr(B)\Pr(W_b/B)}{\Pr(B)\Pr(W_b/B) + \Pr(G)\Pr(W_b/G)}$$

$$\Pr(B/W_b) = (.15 \times .8) \;/\; [(.15 \times .8) + (.85 \times .2)] = 12/29 \approx 0.41$$

Answer:

$\Pr(B/W_b) \approx 0.41$.
$\Pr(G/W_b) \approx 1 - 0.41 = 0.59$.
It is more likely that the sideswiper was green.

BASE RATES

Why do so few people feel, intuitively, that (d) is the right answer? Tversky and Kahneman argue that people tend to ignore the *base rate* or background information. We focus on the fact that the witness is right 80% of the time. We ignore the fact that most of the cabs in town are green.

Suppose that we made a great many experiments with the witness, randomly selecting cabs and showing them to her on a misty night. If 100 cabs were picked at random, then we'd expect something like this:

> The witness sees about 85 green cabs. She correctly identifies 80% of these as green: about 68.
> She incorrectly identifies 20% as blue: about 17.
> She sees about 15 blue cabs. She correctly identifies 80% of these as blue: about 12.
> She incorrectly identifies 20% as green: about 3.

So the witness identifies about 29 cabs as blue, but only 12 of these are blue! In fact, the more we think of the problem as one about frequencies, the clearer the Bayesian answer becomes.

Some critics say that the taxicab problem does not show that we make mistakes easily. The question is asked in the wrong way. If we had been asked just about frequencies, say the critics, we would have given pretty much the right answer straightaway!

RELIABILITY

Our witness was pretty reliable: right 80% of the time. How can a reliable witness not be trustworthy? Because of the base rates. We tend to confuse two different ideas of "reliability."

Idea 1: $\Pr(W_b/B)$: How reliable is she at identifying a cab as blue, given that it is in fact blue? This is a characteristic of the witness and her perceptual acumen.

Idea 2: $\Pr(B/W_b)$: How well can what the witness said be relied on, given that she said the cab is blue? This is a characteristic of the witness and the base rate.

FALSE POSITIVES

Base rates are very striking with medical diagnoses. Suppose I am tested for a terrible disease. I am told that the test is 99% right. If I have the disease, the test says YES with probability 99%. If I do not have the disease, it says NO with probability 99%.

I am tested for the disease. The test says YES. I am terrified.

But suppose the disease is very rare. In the general population, only one person in 10,000 has this disease.

Then among one million people, only 100 have the disease.

In testing a million people at random, our excellent test will answer YES for about 1% of the population, that is, 10,000 people. But as we see by a simple calculation in the next section, *at most 100 of these people actually have the disease*! I am relieved, unless I am in a population especially at risk.

I was terrified by a result YES, plus the test "reliability" (*Idea 1*):

Pr(YES/I'm sick).

But I am relieved once I find out about the "reliability" of a test result (*Idea 2*):

Pr(I'm sick/YES).

A test result of YES, when the correct answer is NO, is called a *false positive*. In our example, about 9,900 of the YES results were false positives.

Thus even a very "reliable" test may be quite misleading, if the base rate for the disease is very low. Exactly this argument was used against universal testing for the HIV virus in the entire population. Even a quite reliable test would give far too many false positives. Even a reliable test can be trusted only when applied to a population "at risk," that is, where the base rate for the disease is substantial.

PROBABILITY OF A FALSE POSITIVE

The result of testing an individual for a condition D is *positive* when according to the test the individual has the condition D.

The result of testing an individual for a condition D is a *false positive* when the individual does not have condition D, and yet the test result is nevertheless positive.

How much can we rely on a test result? This is *Idea 2* about reliability. The probability of a false positive is a good indicator of the extent to which you should rely on (or doubt) a test result.

Let D be the hypothesis that an individual has condition D.
Let Y be YES, a positive test result for an individual.
A false positive occurs when an individual does not have condition D, even though the test result is Y.
The probability of a false positive is Pr(~D/Y).

In our example of the rare disease:

The base rate is Pr(D) = 1/10,000. Hence Pr(~D) = 9,999/10,000.
The test's "reliability" (*Idea 1*) is Pr(Y/D) = 0.99.
And Pr(Y/~D) = 0.01.

Applying Bayes' Rule,

$$\Pr(-D/Y) = \frac{\Pr(-D)\Pr(Y/-D)}{\Pr(\sim D)\Pr(Y/\sim D) + \Pr(D)\Pr(Y/D)} = 9999/(9999 + 99) \approx 0.99.$$

STREP THROAT: ODD QUESTION 6

6. You are a physician. You think it is quite likely that one of your patients has strep throat, but you aren't sure. You take some swabs from the throat and send them to a lab for testing. The test is (like nearly all lab tests) not perfect.

   If the patient has strep throat, then 70% of the time the lab says YES. But 30% of the time it says NO.
   If the patient does not have strep throat, then 90% of the time the lab says NO. But 10% of the time it says YES.

   You send five successive swabs to the lab, from the same patient. You get back these results, in order:

   YES, NO, YES, NO, YES

   You conclude:
   ______(a) These results are worthless.
   ______(b) It is likely that the patient does **not** have strep throat.
   ______(c) It is **slightly** more likely than not, that the patient **does** have strep throat.
   ______(d) It is **very much more** likely than not, that the patient **does** have strep throat.

In my experience almost no one finds the correct answer very obvious. It looks as if the yes-no-yes-no-yes does not add up to much. In fact, it is very good evidence that your patient has strep throat.

Let S = the patient has strep throat.
Let ~S = the patient does not have strep throat.
Let Y = a test result is positive.
Let N = a test result is negative

You think it likely that the patient has strep throat. Let us, to get a sense of the problem, put a number to this, a probability of 90%, that the patient has strep throat. Pr(S) = 0.9.

Solution

We know the conditional probabilities, and we assume that test outcomes are independent.

$\Pr(Y/S) = 0.7 \qquad \Pr(N/S) = 0.3$
$\Pr(Y/{\sim}S) = 0.1 \qquad \Pr(N/{\sim}S) = 0.9$

We need to find Pr(S/YNYNY).

$\Pr(YNYNY/S) = 0.7 \times 0.3 \times 0.7 \times 0.3 \times 0.7 = 0.03087$
$\Pr(YNYNY/{\sim}S) = 0.1 \times 0.9 \times 0.1 \times 0.9 \times 0.1 = 0.00081$

$$\Pr(S/YNYNY) = \frac{\Pr(S)\Pr(YNYNY/S)}{\Pr(S)\Pr(YNYNY/S) + \Pr({\sim}S)\Pr(YNYNY/{\sim}S)}$$

$$\Pr(S/YNYNY) = \frac{0.9 \times 0.03087}{(0.9 \times 0.03087) + (0.1 \times 0.00081)} = 0.997$$

Or you can do the calculation with the original figures, most of which cancel, to give Pr(S/YNYNY) = 343/344. Starting with a prior assumption that Pr(S) = 0.9, we have found that Pr(S/YNYNY) is almost 1!

Answer: So (d) is correct: *It is very much more* likely than not, that the patient *does* have strep throat.

SHEER IGNORANCE

But you are not a physician. You cannot read the signs well. You might just as well toss a coin to decide whether your friend has strep throat. You would model your ignorance as tossing a coin:

$\Pr(S) = 0.5.$

Then you learn of the test results. Should they impress you, or are they meaningless? You require Pr(S/YNYNY).

Solution

Using the same formula as before, but with Pr(S) =0.5,
$\Pr(S/YNYNY) = (.5 \times .03087)/[(.5 \times .03087) + (.5 \times .00081)] \approx 0.974.$
Or, exactly, 343/352.

Answer: This result shows once again that the test results YNYNY are *powerful* evidence that your friend has strep throat.

REV. THOMAS BAYES

Bayes' Rule is named after Thomas Bayes (1702–1761), an English minister who was interested in probability and induction. He probably disagreed strongly with the Scottish philosopher David Hume about evidence. Chapter 21 explains how one might evade Hume's philosophical problem about induction by using Bayesian ideas.

Bayes wrote an essay that was published in 1763 (after his death). It contains the solution to a sophisticated problem like the examples given above. He imagines that a ball is thrown onto a billiard table. The table is "so made and leveled" that a ball is as likely to land on any spot as on any other. A line is drawn through the ball, parallel to the ends of the table. This divides the table into two parts, A and B, with A at a distance of a inches from one end.

Now suppose you do not know the value of a. The ball has been thrown behind your back, and removed by another player.

Then the ball is thrown n times. You are told that on k tosses the ball falls in segment A of the table, and in $n-k$ tosses it falls in segment B. Can you make a guess, on the basis of this information, about the value of a? Obviously, if most of the balls fell in A, then a must cover most of the length of the table; if it is about 50:50 A and B, then a should be about half the length of the table.

Thomas Bayes shows how to solve this problem exactly, finding, for any distance x, and any interval ε, the probability that the unknown a lies between $(x-\varepsilon)$ and $(x+\varepsilon)$.

The idea he used is the same as in our examples, but the mathematics is hard. What is now called *Bayes' Rule* (or, misleadingly, Bayes' Theorem) is a trivial simplification of Bayes' work. In fact, as we saw in Chapter 4, all the work we do with Bayes' Rule can be done from first principles, starting with the definition of conditional probability.

EXERCISES

1 *Lamps and triangles.* Use Bayes' Rule to solve 2(c), and 3(c) in the exercises for Chapter 5, page 56.

2 *Double dipping.*

Contents of urn A: 60 red, 40 green balls.
Contents of urn B: 10 red, 90 green balls.

An urn is chosen by flipping a fair coin.
(a) Two balls are drawn from this urn with replacement. Both are red. What is the probability that we have urn A?
(b) Two balls are drawn from this urn without replacement. Both are red. What is the probability that we have urn A?

3 *Tests.* A professor gives a true-false examination consisting of thirty T-F questions. The questions whose answers are "true" are randomly distributed among the thirty questions. The professor thinks that ¾ of the class are serious, and have correctly mastered the material, and that the probability of a correct answer on any question from such students is 75%. The remaining students will answer at random. She glances at a couple of questions from a test picked haphazardly. Both questions are answered correctly. What is the probability that this is the test of a serious student?

4 *Weightlifters.* Recall the coach that sent one of two teams for competition (page 54 above). Each team has ten members. Eight members of the Steroid team (S) use steroids (U). Two members of the Cleaner team (C) use steroids. The coach chooses which team to send for competition by tossing a fair coin.

One athletics committee tests for steroids in the urine of only one randomly chosen member of the team that has been sent. The test is 100% effective. If this team member is a user, the team is rejected.

(a) What would be a false positive rejection of the entire team?

(b) What is the probability of a false positive?

(c) Another committee is more rigorous. It randomly chooses two different members. What is the probability of a false positive?

5 *Three hypotheses.* (a) State Bayes' Rule for the conditional probability Pr(F/E) with three mutually exclusive and exhaustive hypotheses, F, G, H. (b) Prove it.

6 *Computer crashes.* A small company has just bought three software packages to solve an accounting problem. They are called Fog, Golem, and Hotshot. On first trials, Fog crashes 10% of the time, Golem 20% of the time, and Hotshot 30% of the time.

Of ten employees, six are assigned Fog, three are assigned Golem, and one is assigned Hotshot. Sophia was assigned a program at random. It crashed on the first trial. What is the probability that she was assigned Hotshot?

7 *Deterring burglars.* This example is based on a letter that a sociologist wrote to the daily newspaper. He thinks that it is a good idea for people to have handguns at home, in order to deter burglars. He states the following (amazing) information:

> The rate with which a home in the United States is burgled at least once per year is 10%. The rate for Canada is 40%, and for Great Britain is 60%. These rates have been stable for the past decade.

Don't believe everything a professor says, especially when he writes to the newspaper! Suppose, however, that the information is correct as stated, and that:

> Jenny Park, Larry Chen, and Ali Sami were trainee investment bankers for a multinational company. During the last calendar year Jenny had a home in the United States, Larry in Great Britain, and Ali in Canada.

One of the trainees is picked at random. This person was burgled last year. What is the probability that this person was Ali?

KEY WORDS FOR REVIEW

Bayes' Rule
Base rates
False positives