

Howson 2

Bayes's Theorem

(pp. 13–26)

Patrick Maher
Philosophy 471

Fall 2006

Notation

symbol	meaning
a, b, c, \dots	any propositions
$\sim a$	not- a
$a \vee b$	a or b
$a \& b$	a and b
$P(a)$	the probability of a
$P(a b)$	the probability of a given b

Axioms of probability (p. 16)

- 1 $P(a) \geq 0$, for all a .
- 2 $P(a) = 1$ if a is logically true.
- 3 $P(a \vee b) = P(a) + P(b)$ if $a \& b$ is logically false.
- 4 $P(a|b) = \frac{P(a \& b)}{P(b)}$ if $P(b) \neq 0$.

Bayes's theorem (p. 21)

If $P(h_1 \vee \dots \vee h_n) = 1$ and $P(h_i \& h_j) = 0$ for $i \neq j$ then

$$P(h_i|e) = \frac{P(e|h_i) P(h_i)}{P(e|h_1) P(h_1) + \dots + P(e|h_n) P(h_n)}.$$

In typical applications, the h_i are hypotheses and e is observational evidence. Terminology:

- $P(h_i)$: The *prior probability* of h_i .
- $P(h_i|e)$: The *posterior probability* of h_i .
- $p(e|h_i)$: The *likelihood* of h_i .

Example

An urn contains 10 dice, of which 9 are fair but 1 is biased so it lands six half the time. You randomly draw a die from the urn and toss it; it lands six. What is the probability that you drew the biased die?

- e = die lands six.
- h_1 = you drew the biased die, h_2 = you drew a fair die.

$$\begin{aligned}P(h_1|e) &= \frac{P(e|h_1) P(h_1)}{P(e|h_1) P(h_1) + P(e|h_2) P(h_2)} \\&= \frac{(1/2)(1/10)}{(1/2)(1/10) + (1/6)(9/10)} \\&= 1/4.\end{aligned}$$

The following problem was given to students and staff at Harvard Medical School.

The Harvard Medical School Test (p. 22)

A diagnostic test for a disease, D , has two outcomes “positive” and “negative” (supposedly indicating the presence and absence of D respectively). The test is a fairly sensitive one: its chance of giving a false negative outcome (showing “negative” when the subject has D) is equal to 0, and its chance of giving a false positive outcome (showing “positive” when the subject does not have D) is small: let us suppose it is equal to 5%. Suppose the incidence of the disease is very low, say one in one thousand in the population. A randomly selected person is given the test and shows a positive outcome. What is the chance they have D ?

- The majority of respondents said the chance is 95%.
- In fact, it is less than 2%, as we can show using Bayes's theorem.

Application of Bayes's theorem to the Harvard test

- e = the test outcome was positive.
- h = the subject has D .

$$\begin{aligned}P(h|e) &= \frac{P(e|h) P(h)}{P(e|h) P(h) + P(e|\sim h) P(\sim h)} \\ &= \frac{(1)(.001)}{(1)(.001) + (.05)(.999)} \\ &= 0.0196.\end{aligned}$$

Low false-positive and false-negative rates by themselves tell you nothing about how reliable a positive outcome is in any given case: an additional piece of information is required, namely the incidence of the disease in the population. The background incidence also goes by the name of “the base rate,” and thinking that valid inferences can be drawn just from the knowledge of false positive and negative rates has come to be called the “base-rate fallacy.”

Significance tests

Standard statistical tests are meant to determine whether some hypothesis h is true. The statistician conducts an experiment and adopts a rule of this form:

- If e is observed, accept h .
- If $\sim e$ is observed, reject h .

Statisticians say a test is good if these are both small:

- $P(\sim e|h)$: The probability of rejecting h when it is true; this is called the *significance level* of the test.
- $p(e|\sim h)$: The probability of accepting h when it is false.

But these probabilities . . . are, in effect, just the chances of a false negative and a false positive, and as we saw so graphically in the Harvard Medical School Test, finding an outcome in such a region conveys no information whatever by itself about the chance of the hypothesis under test being true. (p. 26)

The No Miracles Argument for scientific realism

Scientific realism is the view that our best scientific theories are at least approximately true. A popular argument for it:

- If our best theories were not at least approximately true, then their agreement with the data would be a miracle.
- We should avoid postulating miracles.
- Therefore, our best scientific theories are at least approximately true.

In probabilistic terms: If h is a successful scientific theory, and e the data that agrees with it, $P(e|\sim h)$ is extremely small, whereas $P(e|h)$ is high; therefore, $P(h|e)$ is high.

Again, we see essentially the same fallacious inference based on a small false positive rate and a small false negative rate as was committed by the respondents to the Harvard Test. (p. 26)

- 1 State Bayes's theorem (not just the formula, but also the conditions under which it holds).
- 2 In a certain town, 85% of the cabs are green and the other 15% are blue. One night a cab is involved in a hit-and-run accident and a witness says the cab was blue. The witness was tested under conditions like those on the night of the accident and found to correctly identify blue cabs as blue 80% of the time, and likewise for green cabs. On the basis of this information, what is the probability that the cab was blue?

- ③ A statistics textbook says:

A good test procedure is one in which both α [the probability of rejecting the hypothesis when it is true] and β [the probability of accepting it when it is false] are small, thereby giving us a good chance of making the correct decision.

(Freund and Walpole, *Mathematical Statistics*, 3rd ed. p. 364.)

What is wrong with this? Justify your answer.

- ④ John Worrall wrote:

*It would be a miracle, a coincidence on a near-cosmic scale, if a theory made as many correct empirical predictions as, say, the general theory of relativity or the photon theory of light without what the theory says about the fundamental structure of the universe being correct or “essentially” or “basically” correct. But we shouldn’t accept miracles, not at any rate if there is a non-miraculous alternative . . . So it is plausible to conclude that presently accepted theories are indeed “essentially” correct. (Quoted in Howson, *Hume’s Problem*, p. 37)*

Is this correct? Explain.