

Prediction and Odds

18.05 Spring 2017



Probabilistic Prediction

Also called **probabilistic forecasting**.

Assign a probability to each outcome of a **future experiment**.

Prediction: “It will rain tomorrow.”

Probabilistic prediction: “Tomorrow it will rain with probability 60% (and not rain with probability 40%).”

Examples: medical treatment outcomes, weather forecasting, climate change, sports betting, elections, ...

Words of estimative probability (WEP)

WEP Prediction: “It is likely to rain tomorrow.”

Memo: *Bin Laden Determined to Strike in US*

See http://en.wikipedia.org/wiki/Words_of_Estimative_Probability

“The language used in the [Bin Laden] memo lacks words of estimative probability (WEP) that reduce uncertainty, thus preventing the President and his decision makers from implementing measures directed at stopping al Qaeda’s actions.”

“Intelligence analysts would rather use words than numbers to describe how confident we are in our analysis,” a senior CIA officer who’s served for more than 20 years told me. Moreover, “most consumers of intelligence aren’t particularly sophisticated when it comes to probabilistic analysis. They like words and pictures, too. My experience is that [they] prefer briefings that don’t center on numerical calculation.”

WEP versus Probabilities: medical consent

No common standard for converting WEP to numbers.

Suggestion for potential risks of a medical procedure:

Word	Probability
Likely	Will happen to more than 50% of patients
Frequent	Will happen to 10-50% of patients
Occasional	Will happen to 1-10% of patients
Rare	Will happen to less than 1% of patients

From same Wikipedia article

Example: Three types of coins

- Type A coins are fair, with probability 0.5 of heads
- Type B coins have probability 0.6 of heads
- Type C coins have probability 0.9 of heads

A drawer contains one coin of each type. You pick one at random.

Prior predictive probability: Before taking data, what is the probability a toss will land heads? Tails?

Take data: say the first toss lands heads.

Posterior predictive probability: After taking data. What is the probability the next toss lands heads? Tails?

Solution 1

1. Use the law of total probability: (A probability tree is an excellent way to visualize this. You should draw one before reading on.)

Let $D_{1,H}$ = 'toss 1 is heads', $D_{1,T}$ = 'toss 1 is tails'.

$$\begin{aligned}P(D_{1,H}) &= P(D_{1,H}|A)P(A) + P(D_{1,H}|B)P(B) + P(D_{1,H}|C)P(C) \\&= 0.5 \cdot 0.3333 + 0.6 \cdot 0.3333 + 0.9 \cdot 0.3333 \\&= 0.6667\end{aligned}$$

$$P(D_{1,T}) = 1 - P(D_{1,H}) = 0.3333$$

Solution 2

2. We are given the data $D_{1,H}$. First update the probabilities for the type of coin.

Let $D_{2,H}$ = 'toss 2 is heads', $D_{2,T}$ = 'toss 2 is tails'.

hypothesis	prior	likelihood	Bayes	
			numerator	posterior
H	$P(H)$	$P(D_{1,H} H)$	$P(D_{1,H} H)P(H)$	$P(H D_{1,H})$
A	1/3	0.5	0.1667	0.25
B	1/3	0.6	0.2	0.3
C	1/3	0.9	0.3	0.45
total	1		0.6667	1

Next use the law of total probability:

$$\begin{aligned}P(D_{2,H}|D_{1,H}) &= P(D_{2,H}|A)P(A|D_{1,H}) + P(D_{2,H}|B)P(B|D_{1,H}) \\&\quad + P(D_{2,H}|C)P(C|D_{1,H}) \\&= 0.71\end{aligned}$$

$$P(D_{2,T}|D_{1,H}) = 0.29.$$

Three coins, continued.

As before: 3 coins with probabilities 0.5, 0.6, and 0.9 of heads.

Pick one; toss 5 times; Suppose you get 1 head out of 5 tosses.

Concept question: What's your best guess for the probability of heads on the next toss?

(a) 0.1

(b) 0.2

(c) 0.3

(d) 0.4

(e) 0.5

(f) 0.6

(g) 0.7

(h) 0.8

(i) 0.9

(j) 1.0

Three coins: a fountain of possibilities.

As before: 3 coins with probabilities 0.5, 0.6, and 0.9 of heads.

Pick one; toss 5 times; Suppose you get 4 heads out of 5 tosses.

Concept question: What's your best guess for the probability of heads on the next toss?

- | | | | |
|---------|---------|---------|---------|
| (0) 0.0 | (1) 0.1 | (2) 0.2 | (3) 0.3 |
| (4) 0.4 | (5) 0.5 | (6) 0.6 | (7) 0.7 |
| (8) 0.8 | (9) 0.9 | | |

Board question: three coins

Same setup:

- 3 coins with probabilities 0.5, 0.6, and 0.9 of heads.
- Pick one; toss 5 times.
- Suppose you get 1 head out of 5 tosses.

Compute the posterior probabilities for the type of coin and the posterior predictive probabilities for the results of the next toss.

1. Specify clearly the set of hypotheses and the prior probabilities.
2. Compute the prior and posterior predictive distributions, i.e. give the probabilities of all possible outcomes.

answer: *See next slide.*

Solution

Data = '1 head and 4 tails'

hypothesis	prior	likelihood	Bayes numerator	posterior
H	$P(H)$	$P(D H)$	$P(D H)P(H)$	$P(H D)$
A	$1/3$	$\binom{5}{1}(0.5)^5$	0.0521	0.669
B	$1/3$	$\binom{5}{1}(0.6) \cdot (0.4)^4$	0.0256	0.329
C	$1/3$	$\binom{5}{1}(0.9) \cdot (0.1)^4$	0.00015	0.002
total	1		0.0778	1

So,

$$P(\text{heads}|D) = 0.669 \cdot 0.5 + 0.329 \cdot 0.6 + 0.002 \cdot 0.9 = 0.53366$$

$$P(\text{tails}|D) = 1 - P(\text{heads}|D) = 0.46634.$$

Concept Question

Does the order of the 1 head and 4 tails affect the posterior distribution of the coin type?

1. Yes
2. No

Does the order of the 1 head and 4 tails affect the posterior predictive distribution of the next flip?

1. Yes
2. No

answer: No for both questions.

Odds

Definition The **odds** of an event are

$$O(E) = \frac{P(E)}{P(E^c)}.$$

- Usually for two choices: E and *not* E .
- Can split multiple outcomes into two groups.
- Can do odds of A vs. $B = P(A)/P(B)$.
- Our Bayesian focus:
Updating the odds of a hypothesis H given data D .
- **posterior odds = prior odds times likelihood ratio**

This simple formula can be a reason to prefer odds to probabilities.

Examples

- A fair coin has $O(\text{heads}) = \frac{0.5}{0.5} = 1$.

We say '1 to 1' or 'fifty-fifty'.

- The odds of rolling a 4 with a six-sided die are $\frac{1/6}{5/6} = \frac{1}{5}$.

We say '1 to 5 for' or '5 to 1 against'

- For event E , if $P(E) = p$ then $O(E) = \frac{p}{1-p}$.

- If an event is **rare**, then $P(E) \approx O(E)$.

Bayesian framework: Marfan's Syndrome

Marfan's syndrome (M) is a genetic disease of connective tissue. The main ocular features (F) of Marfan syndrome include bilateral ectopia lentis (lens dislocation), myopia and retinal detachment.

$$P(M) = 1/15000, \quad P(F|M) = 0.7, \quad P(F|M^c) = 0.07$$

If a person has the main ocular features F what is the probability they have Marfan's syndrome?

hypothesis	prior	likelihood	Bayes	
			numerator	posterior
\mathcal{H}	$P(\mathcal{H})$	$P(F \mathcal{H})$	$P(F \mathcal{H})P(\mathcal{H})$	$P(\mathcal{H} F)$
M	0.000067	0.7	0.0000467	0.00066
M^c	0.999933	0.07	0.069995	0.99933
total	1		0.07004	1

Odds form

$$P(M) = 1/15000, \quad P(F|M) = 0.7, \quad P(F|M^c) = 0.07$$

Prior odds:

$$O(M) = \frac{P(M)}{P(M^c)} = \frac{1/15000}{14999/15000} = \frac{1}{14999} = 0.000067.$$

Note: $O(M) \approx P(M)$ since $P(M)$ is small.

Posterior odds: can use the Bayes numerator!

$$O(M|F) = \frac{P(M|F)}{P(M^c|F)} = \frac{P(F|M)P(M)}{P(F|M^c)P(M^c)} = 0.000667.$$

The posterior odds is a product of factors:

$$O(M|F) = \frac{P(F|M)}{P(F|M^c)} \cdot \frac{P(M)}{P(M^c)} = \frac{0.7}{0.07} \cdot O(M)$$

Bayes factors

$$\begin{aligned}O(M|F) &= \frac{P(F|M)}{P(F|M^c)} \cdot \frac{P(M)}{P(M^c)} \\&= \frac{P(F|M)}{P(F|M^c)} \cdot O(M)\end{aligned}$$

posterior odds = Bayes factor \cdot prior odds

- The Bayes factor is the ratio of the likelihoods.
- The Bayes factor gives the strength of the ‘evidence’ provided by the data.
- A large Bayes factor times small prior odds can be small (or large or in between).
- The Bayes factor for ocular features is $0.7/0.07 = 10$.

Board Question: screening tests

A disease is present in 0.005 of the population.

A screening test has a 0.05 false positive rate and a 0.02 false negative rate.

1. Give the prior odds a patient has the disease

Assume the patient tests positive

2. What is the Bayes factor for this data?

3. What are the posterior odds they have the disease?

4. Based on your answers to (1) and (2) would you say a positive test (the data) provides strong or weak evidence for the presence of the disease.

answer: See next slide

Solution

Let \mathcal{H}_+ = 'has disease' and \mathcal{H}_- = 'doesn't'

Let \mathcal{T}_+ = positive test

$$1. O(\mathcal{H}_+) = \frac{P(\mathcal{H}_+)}{P(\mathcal{H}_-)} = \frac{0.005}{0.995} = 0.00503$$

Likelihood table:

	Possible data	
	\mathcal{T}_+	\mathcal{T}_-
Hypotheses $\left\{ \begin{array}{l} \mathcal{H}_+ \\ \mathcal{H}_- \end{array} \right.$	0.98	0.02
	0.05	0.95

$$2. \text{ Bayes factor} = \text{ratio of likelihoods} = \frac{P(\mathcal{T}_+|\mathcal{H}_+)}{P(\mathcal{T}_+|\mathcal{H}_-)} = \frac{0.98}{0.05} = 19.6$$

$$3. \text{ Posterior odds} = \text{Bayes factor} \times \text{prior odds} = 19.6 \times 0.00504 = 0.0985$$

4. Yes, a Bayes factor of 19.6 indicates a positive test is strong evidence the patient has the disease. The posterior odds are still small because the prior odds are **extremely** small.

More on next slide.

Solution continued

Of course we can compute the posterior odds by computing the posterior probabilities using a Bayesian update table.

hypothesis	prior	likelihood	Bayes	
			numerator	posterior
\mathcal{H}	$P(\mathcal{H})$	$P(\mathcal{T}_+ \mathcal{H})$	$P(\mathcal{T}_+ \mathcal{H})P(\mathcal{H})$	$P(\mathcal{H} \mathcal{T}_+)$
\mathcal{H}_+	0.005	0.98	0.00490	0.0897
\mathcal{H}_-	0.995	0.05	0.04975	0.9103
total	1		0.05474	1

Posterior odds: $O(\mathcal{H}_+|\mathcal{T}_+) = \frac{0.0897}{0.9103} = 0.0985$

Board Question: CSI Blood Types*

- Crime scene: the two perpetrators left blood: one of type O and one of type AB
- In population 60% are type O and 1% are type AB
- ① Suspect Oliver is tested and has type O blood.
Compute the Bayes factor and posterior odds that Oliver was one of the perpetrators.

Is the data evidence for or against the hypothesis that Oliver is guilty?
- ② Same question for suspect Alberto who has type AB blood.

Show helpful hint on next slide.

*From 'Information Theory, Inference, and Learning Algorithms' by David J. C. Mackay.

Helpful hint

Population: 60% type O; 1% type AB

For the question about Oliver we have

Hypotheses:

S = 'Oliver and another unknown person were at the scene'

S^c = 'two unknown people were at the scene'

Data:

D = 'type 'O' and 'AB' blood were found; Oliver is type O'

Solution to CSI Blood Types

For Oliver:

$$\text{Bayes factor} = \frac{P(D|S)}{P(D|S^c)} = \frac{0.01}{2 \cdot 0.6 \cdot 0.01} = 0.83.$$

Therefore the posterior odds $= 0.83 \times \text{prior odds}$ ($O(S|D) = 0.83 \cdot O(S)$)
Since the odds of his presence decreased this is (weak) evidence of his innocence.

For Alberto:

$$\text{Bayes factor} = \frac{P(D|S)}{P(D|S^c)} = \frac{0.6}{2 \cdot 0.6 \cdot 0.01} = 50.$$

Therefore the posterior odds $= 50 \times \text{prior odds}$ ($O(S|D) = 50 \cdot O(S)$)
Since the odds of his presence increased this is (strong) evidence of his presence at the scene.

David Mackay:

“In my view, a jury’s task should generally be to multiply together carefully evaluated likelihood ratios from each independent piece of admissible evidence with an equally carefully reasoned prior probability. This view is shared by many statisticians but learned British appeal judges recently disagreed and actually overturned the verdict of a trial because the jurors *had* been taught to use Bayes’ theorem to handle complicated DNA evidence.”

Updating again and again

Collect data: D_1, D_2, \dots

Posterior odds to D_1 become prior odds to D_2 . So,

$$\begin{aligned} O(H|D_1, D_2) &= O(H) \cdot \frac{P(D_1|H)}{P(D_1|H^c)} \cdot \frac{P(D_2|H)}{P(D_2|H^c)} \\ &= O(H) \cdot BF_1 \cdot BF_2. \end{aligned}$$

Independence assumption:

D_1 and D_2 are *conditionally* independent.

$$P(D_1, D_2|H) = P(D_1|H)P(D_2|H).$$

Marfan's Symptoms

The Bayes factor for ocular features (F) is

$$BF_F = \frac{P(F|M)}{P(F|M^c)} = \frac{0.7}{0.07} = 10$$

The wrist sign (W) is the ability to wrap one hand around your other wrist to cover your pinky nail with your thumb. Assume 10% of the population have the wrist sign, while 90% of people with Marfan's have it. So,

$$BF_W = \frac{P(W|M)}{P(W|M^c)} = \frac{0.9}{0.1} = 9.$$

$$O(M|F, W) = O(M) \cdot BF_F \cdot BF_W = \frac{1}{14999} \cdot 10 \cdot 9 \approx \frac{6}{1000}.$$

We can convert posterior odds back to probability, but since the odds are so small the result is nearly the same:

$$P(M|F, W) \approx \frac{6}{1000 + 6} \approx 0.596\%$$