More on Decisions and Hypotheses

PUBH 8442: Bayes Decision Theory and Data Analysis

Eric F. Lock
UMN Division of Biostatistics, SPH
elock@umn.edu

02/17/2021

Hypothesis decisions

- Often, no need to conclude or decide on a given hypothesis
  - Simply report strength of evidence given by p-value or posterior
- If necessary to choose, can use decision-theoretic framework
- Action space $\mathcal{A} = \{H_0, H_a\}$
- A natural decision rule is

$$d_c(\mathbf{y}) = \begin{cases} H_0 & \text{if } P(H_0 \mid \mathbf{y}) > c \\ H_a & \text{otherwise} \end{cases}$$

for $c \in (0, 1)$

  - $c = 1/2$ chooses whichever hypothesis is more probable

Type I error

- A *type I error* occurs when H_0 is true, but we conclude it is not.
- If we observe $\mathbf{y}$ and conclude H_a, the probability we made a type I error is $P(H_0 \mid \mathbf{y})$
- The standard definition of type I error rate, for a given rule, is

$$P(d(\mathbf{y}) \neq H_0 \mid H_0)$$

- Type I error rate for rule d_c depends on context.

Example: tipping pennies (cont.)

- Bayesian framework:
  - Let $\theta =$ probability the penny lands heads
  - $H_0 : \theta = 1/2$
  - $H_a : \theta \sim \text{Uniform}(0.5, 1)$
  - $P(H_0) = 0.5$
- Two different experiments, resulting in sampling distributions
  - $X \sim \text{NegativeBinomial}(1, \theta)$
  - $Y \sim \text{Binomial}(6, \theta)$
- Consider $d_{0.3}$: rule choosing H_0 if $P(H_0 \mid \mathbf{y}) > 0.3$

Example: tipping pennies (cont.)

x	$P(H_0 \mid x)$	y	$P(H_0 \mid y)$
0	0.67	0	0.88
1	0.60	1	0.84
2	0.52	2	0.78
3	0.43	3	0.69
4	0.34	4	0.51
5	0.26	5	0.26
6	0.18	6	0.05
7	0.13		
$\vdots$	$\vdots$		

http://www.ericfrazerlock.com/More_on_Decisions_and_Hypothesis_Testing_Rcode1.R

- For experiment 1, $d_{0.3}(x) = H_a$ if $x \geq 5$
- For experiment 2, $d_{0.3}(y) = H_a$ if $y \geq 5$

Example: tipping pennies (cont.)

- For experiment 1, type I error rate is $P(X \geq 5 \mid H_0) = 0.031$

$P(X=5 \mid H_0) + P(X=6 \mid H_0) + \ldots$

$= 0.031$

- For experiment 2, type I error rate is $P(Y \geq 5 \mid H_0) = 0.110$

Example: Vaccine trial

- A tentative vaccine trial for a new COVID variant completes
  - 1 of 50 who received the vaccine were infected $\theta \sim Beta(1,1)$
  - 7 of 50 who received a placebo were infected $\theta_1, \theta_2 \overset{iid}{\sim} Beta(1,1)$
- Conduct second stage trial, with $n = 500$ in each group
- Possible Bayesian framework: $p(y_1 | \theta_1) = Bin(n, \theta_1)$
  - Let θ_1 = probability of vaccinated infection,
    θ_2 = probability of non-vaccinated infection
  - $H_0 : \theta_1 = \theta_2 = \theta \sim \text{Beta}(9, 93)$
  - $H_a : \theta_1 \sim \text{Beta}(2, 50), \theta_2 \sim \text{Beta}(8, 44)$ are independent
  - $P(H_0) = 0.25$
- Observe $y_1 = 36$ infections in vaccinated group,
  $y_2 = 52$ in non-vaccinated group

Example: Vaccine trial

- The observed probability under H_0 is

$$P(y_1 = 36, y_2 = 52 \mid H_0) \approx 0.0002$$

http://www.ericfrazerlock.com/More_on_Decisions_and_Hypothesis_Testing_Rcode2.r

$$= \int_0^1 P(y_1=36 \mid \theta)\, P(y_2=52 \mid \theta) \cdot p(\theta)\, d\theta$$

$$= \binom{500}{36}\binom{500}{52} \cdot \frac{1}{B(9,93)} \int_0^1 \theta^{36}(1-\theta)^{464}\, \theta^{52}(1-\theta)^{448} \cdot \theta^{8}(1-\theta)^{92}\, d\theta$$

$$= \int_0^1 \theta^{96}(1-\theta)^{1004}\, d\theta$$

$$\propto \text{Beta}(97, 1005) = \binom{500}{36}\binom{500}{52} \frac{B(97,1005)}{B(9,93)}$$

- The observed probability under H_a is

$$P(y_1 = 36, y_2 = 52 \mid H_a) \approx 0.0001.$$

$$= \int_0^1 \int_0^1 p(y_1=36 \mid \theta_1) \cdot p(y_2=52 \mid \theta_2) \cdot p(\theta_1 \mid H_a) \cdot p(\theta_2 \mid H_a)\, d\theta_1\, d\theta_2$$

$$= \int_0^1 p(y_1=36 \mid \theta_1) \cdot p(\theta_1 \mid H_a)\, d\theta_1 \cdot \int_0^1 p(y_2=52 \mid \theta_2) \cdot p(\theta_2 \mid H_a)\, d\theta$$

$$= \ldots = \binom{500}{36} \cdot \frac{B(38,514)}{B(2,50)} \cdot \binom{500}{52} \cdot \frac{B(60,492)}{B(8,44)}$$

Example: Vaccine trial

- The posterior probability of H_0 is

$$P(H_0 \mid y_1 = 36, y_2 = 52) \approx 0.42$$

- Our probability that the vaccine has an effect is 0.58.
    - Less sure than before

Example: Vaccine trial

- By the beta-binomial model,
  - Under H_0,

    $p(\theta_1, \theta_2 \mid H_0, y) \neq p(\theta_1 \mid H_0, y) \cdot p(\theta_2 \mid H_0, y)$

    $$p(\theta_1 = \theta_2 = \theta \mid H_0, \mathbf{y}) = \text{Beta}(9 + y_1 + y_2, 93 + 1000 - y_1 - y_2)$$

    $\{y_1, y_2\}$

  - Under H_a,

    $$p(\theta_1 \mid H_a, \mathbf{y}) = \text{Beta}(2 + y_1, 50 + 500 - y_1)$$

    and

    $$p(\theta_2 \mid H_a, \mathbf{y}) = \text{Beta}(8 + y_2, 44 + 500 - y_2)$$

- The marginal distributions over H_0, H_a for θ_1, θ_2 are

$$p(\theta_i \mid \mathbf{y}) = P(H_0 \mid \mathbf{y}) p(\theta \mid H_0, \mathbf{y}) + P(H_a \mid \mathbf{y}) p(\theta_i \mid H_a, \mathbf{y})$$

Example: Vaccine trial

- For $y_1 = 36$ and $y_2 = 52$,
  - Under H_0,

$$p(\theta_1 = \theta_2 = \theta \mid H_0, \mathbf{y}) = \text{Beta}(97, 1005)$$

  - Under H_a,

$$p(\theta_1 \mid H_a, \mathbf{y}) = \text{Beta}(38, 514)$$

    and

$$p(\theta_2 \mid H_a, \mathbf{y}) = \text{Beta}(60, 492)$$

- The marginal distributions over H_0, H_a are

$$p(\theta_1 \mid \mathbf{y}) = 0.42\,\text{Beta}(97, 1005) + 0.58\,\text{Beta}(38, 514)$$

and

$$p(\theta_2 \mid \mathbf{y}) = 0.42\,\text{Beta}(97, 1005) + 0.58\,\text{Beta}(60, 492)$$

Example: Vaccine trial

- Marginal densities with (θ_1) and without (θ_2) vaccine:

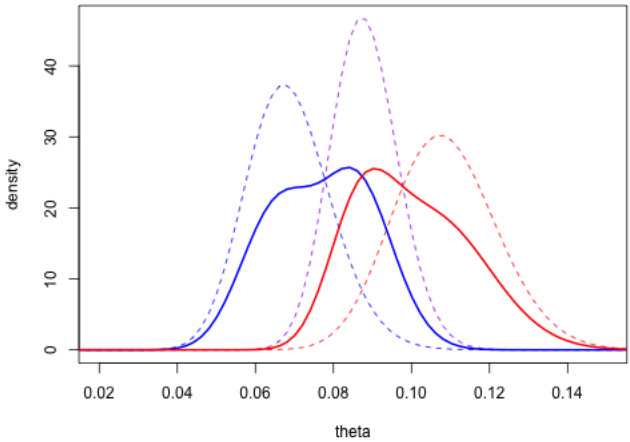

Example: Vaccine trial

- Decide whether to distribute the vaccine, after second trial.
  - Loss for distributing vaccines is $l(\boldsymbol{\theta}, \text{distribute}) = 50000$
  - Loss for not distributing vaccines is

$$l(\boldsymbol{\theta}, \text{do not distribute}) = 10^6(\theta_2 - \theta_1)$$

- Posterior risk for distributing is $\rho(\vec{y}, \text{distribute}) = 50000$
- Posterior risk for not distributing is $E\left[\text{Beta}(a,b)\right] = \frac{a}{a+b}$

$$\rho(\vec{y}, \text{do not distribute}) \approx 23100$$

$E := E_{\theta_1,\theta_2|y}$ $\hookrightarrow$ $E\, 10^6(\theta_2-\theta_1) = 10^6(E\theta_2 - E\theta_1)$

$= 0.42 \cdot \frac{97}{97+1005} + 0.58 \cdot \frac{60}{492+60}$

$- 0.42 \cdot \frac{97}{97+1005} - 0.58 \cdot \frac{38}{38+514}$ $\longrightarrow$

- Choose not to distribute vaccine at this time.

Shrinkage toward H_0

$$\rightarrow 10^6 \cdot \overbrace{0{,}58}\left(\frac{60}{552} - \frac{38}{552}\right)$$

$$\approx 23{,}100$$

Bayes factors

- The posterior *odds* for H_0 over H_a is

$$\frac{P(H_0 \mid \mathbf{y})}{P(H_a \mid \mathbf{y})} = \frac{P(H_0) \cdot P(y \mid H_0)}{P(H_a) \cdot P(y \mid H_a)}$$

- The *Bayes factor* for H_0 over H_a is

$$BF = \frac{p(\mathbf{y} \mid H_0)}{p(\mathbf{y} \mid H_a)}$$

- The Bayes factor is the posterior odds scaled by the prior odds

$$BF = \frac{P(H_a)}{P(H_0)} \cdot \frac{P(H_0 \mid \mathbf{y})}{P(H_a \mid \mathbf{y})}$$

- It is a measure of "objective" evidence for H_0 over H_a

Bayes factors

- Heuristic interpretation of Bayes factors (c. Harold Jeffreys):

BF	Strength of evidence
< 1	Negative
1 to 3	Barely worth mentioning
3 to 10	Substantial
10 to 30	Strong
30 to 100	Very strong
100	Decisive

- Example: The Bayes factor for the second vaccine trial was

$$\frac{P(y_1 = 36, y_2 = 52 \mid H_0)}{P(y_1 = 36, y_2 = 52 \mid H_a)} \approx 2$$

Data gives little evidence for one hypothesis over the other.