

22S:138

The Likelihood Principle

Lecture 21
Dec. 1, 2008

Kate Cowles
374 SH, 335-0727
kcowles@stat.uiowa.edu

Another way to state the likelihood principle

- For a given sample of data, any two probability models $p(\mathbf{y}|\theta)$ that have the same likelihood function yield the same inference for θ .
- With regard to the information contained in the data about the unknown parameter(s), only the actual observed data \mathbf{y} is relevant.
 - No other possible outcomes
 - * Contrast this with the frequentist p-value the probability assuming H_0 is true, of getting a test statistic as extreme as, *or more extreme than*, the value that was actually obtained
 - Not the researchers' intentions

The likelihood principle

- Suppose that two different experiments may inform about an unknown parameter θ
- Suppose the outcomes of the experiments are respectively y^* and z^*
- Suppose the likelihoods for θ resulting from the two experiments are proportional; that is

$$p(y^*; \theta) = c p(z^*; \theta)$$

where c is a constant

- Then the information about θ contained in both experiments is equivalent

Example

- We are given a coin. We are interested in estimating θ , the probability of obtaining a head on a single flip.
- We want to test the hypotheses:

$$H_0 : \theta = \frac{1}{2}$$

$$H_a : \theta > \frac{1}{2}$$
- Experiment consists of flipping coin 12 times independently.
- Result is 9 heads and 3 tails.

Example, continued

- There are (at least) two possible ways the experiment might have been conducted:
 - Design 1: do 12 flips. The random variable Y is the number of heads obtained in $n = 12$ flips.
 - Design 2: Flip the coin until 9 heads are obtained. Random variable Y is the number of tails that are obtained before the ninth head.
- Frequentist inference for θ would be different depending on which design is used.
- Bayesian inference would be the same under both designs because the likelihoods are proportional.
- The negative binomial distribution
 - $Y =$ the number of failures observed in a sequence of independent Bernoulli trials

Implications of the likelihood principle

- the stopping rule principle
- the likelihood principle and reference priors

before the k^{th} success

$$- Y \sim NB(k, p)$$

$$- p(y|p) = \binom{k+y-1}{y} p^k (1-p)^y$$

$$- E(Y) = \frac{k(1-p)}{p}$$

“Stopping rules” are often used in designing frequentist statistical studies

- instead of a fixed sample size
- to make it possible to stop a study early if the results are in
 - particularly common in clinical trials
 - reducing the size and duration of a clinical trial reduces the number of patients who are exposed to the treatment that will be found to be inferior and speeds up the dissemination of the results to the medical community
- Frequentist statisticians must choose the stopping rule *before* the experiment is conducted and adhere to it exactly
 - deviations can produce serious errors if a frequentist analysis is used

- Large frequentist literature on how to control the overall probability of Type I error while allowing for more than one analysis of the data

Stopping Rule Principle

- In a sequential experiment, the evidence provided by the experiment about the value of the unknown parameter(s) θ should not depend on the stopping rule
- follows directly from the likelihood principle

Jeffreys' priors and the likelihood principle

- recall Jeffreys' prior
 - “reference” prior
 - noninformative
 - invariant to transformations of parameters
 - $p(\theta) \propto [I(\theta)]^{-\frac{1}{2}}$
where $I(\theta)$ is the expected Fisher information for θ
- Jeffreys' prior when likelihood is Binomial(n, θ)

$$\begin{aligned} p(\theta) &\propto \theta^{-\frac{1}{2}}(1 - \theta)^{-\frac{1}{2}} \\ &= \text{Beta}\left(\frac{1}{2}, \frac{1}{2}\right) \end{aligned}$$

- Jeffreys' prior when likelihood is negative binomial(k, θ)

$$\begin{aligned} p(\theta) &\propto \theta^{-1}(1 - \theta)^{-\frac{1}{2}} \\ &= \text{Beta}\left(0, \frac{1}{2}\right) \end{aligned}$$

- So use of Jeffreys' prior in some cases can violate the likelihood principle!