STAT 703/J703
April 26th, 2007
-Lecture 28-

Instructor: Brian Habing
Department of Statistics
LeConte 203
Telephone: 803-777-3578
E-mail: habing@stat.sc.edu

Homework Q1:

Find the minimum variance that any unbiased estimator of lambda for an exponential distribution can have.

Chapter 8 #75:

Show that the gamma distribution belongs to the exponential family. Identify the natural parameters and sufficient statistic.

Chapter 9 #36:

Is there any evidence that the suicide rates varies seasonally, or are the data consistent with the hypothesis that this rate is a constant.

Bayesian Statistics

Classical Setup

θ is an unknown constant

$X_1,\ldots,X_n$ are observed from $f(x|\ \underline{\theta})$

Bayesian Setup

Θ is an unknown random variable , unobservable with pdf $g(\underline{\theta})$

$X_1,\ldots,X_n$ are observed from the conditional pdf $f(x|\ \underline{\theta})$

A Bayesian Statistical Model:

1. An observed sample $X_1,\ldots,X_n$, $(X_1,\ldots,X_n)=\underline{X}\in \mathcal{X}$, the data space.
2. An unobserved random vector $\underline{\Theta}$ of parameters, $\underline{\Theta}\in\Omega$, the parameter space.
3. The conditional density of $X_1,\ldots,X_n$, given $\underline{\Theta}=\underline{\theta}$, $f(\underline{x}|\ \underline{\theta})$.
4. The (marginal) pdf of $\underline{\Theta}$, $g(\underline{\theta})$, the prior distribution of $\underline{\Theta}$.

Bayes Rule can be written

$$h(\theta \mid x) = \frac{f(x \mid \theta) g(\theta)}{\int f(x \mid \theta) g(\theta) d\theta}$$

Say we knew $f(x|\theta)$ and $g(\theta)$ and wanted to find the maximum likelihood estimate of $h(\theta|x)$. Why can we just find the value of θ that maximizes $f(x|\theta)\, g(\theta)$ and not worry about the bottom integral?

Hence, inferences are described in terms of the <u>posterior</u> <u>distribution</u> <u>of</u> $\underline{\Theta}$, the conditional pdf of $\underline{\Theta}$, given $\underline{X}=(x_1,\ldots,x_n)$, $h(\underline{\theta}|\underline{x})$. (Interpretation of posterior).

We consider Bayesian <u>point</u> and <u>interval</u> estimators as well as Bayesian <u>tests</u> (<u>decisions</u>). The point estimate is given by a <u>Bayes</u> <u>rule</u>.

<u>Example</u> - Sample $X_1,\ldots,X_n$ from a normal population whose variance is 1 and mean μ is a value of a random variable M,

$X_i|M=\mu \sim N(\mu,1)$,

$$f(x \mid \mu) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2}}, \mu \in R$$

Suppose $M \sim N(\alpha,\tau^2)$ as the prior pdf.

The "Bayesian confidence interval" is called a credibility interval:

(θ_0, θ_1) where $\int_{\theta_1}^{\theta_2} h(\theta \mid x)\, d\theta = 1 - \alpha$

"Bayesian hypothesis testing" depends on a loss function evaluated for competing hypotheses.