557: MATHEMATICAL STATISTICS II
INTERVAL ESTIMATION: WORKED EXAMPLES

Example 1 : Inverting a Test Statistic

Suppose that $X_1, \ldots, X_n \sim \text{Normal}(\theta, \sigma^2)$ for σ^2 known. A confidence interval can be constructed by recalling the UMP unbiased test at level α of

$$\begin{aligned} H_0 &: \theta = \theta_0 \\ H_1 &: \theta \neq \theta_0 \end{aligned}$$

with rejection region

$$\mathcal{R}(\theta_0) \equiv \{\mathbf{x} : \overline{x} < -c_n(\theta_0)\} \cup \{\mathbf{x} : \overline{x} > c_n(\theta_0)\}$$

where

$$c_n(\theta_0) = \frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}} + \theta_0.$$

The corresponding acceptance region is

$$\mathcal{A}(\theta_0) \equiv \{\mathbf{x} : -c_n(\theta_0) < \overline{x} < c_n(\theta_0)\}$$

so that

$$\Pr[\mathbf{X} \in \mathcal{A}(\theta_0); \theta_0] = \Pr[-c_n(\theta_0) < \overline{X} < c_n(\theta_0); \theta_0] = 1 - \alpha.$$

From this we conclude that, under the distribution $f_{\mathbf{X}}(\mathbf{x}; \theta_0)$, we have that the probability that

$$-\frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}} - \theta_0 < \overline{X} < \frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}} + \theta_0$$

is $1-\alpha$. Rearranging, we have that with probability $1-\alpha$,

$$\overline{X} - \frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}} < \theta_0 < \overline{X} + \frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}}.$$

Therefore a $1-\alpha$ confidence interval is defined by $[L(\mathbf{X}), U(\mathbf{X})]$ where

$$L(\mathbf{X}) = \overline{X} - \frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}} \qquad U(\mathbf{X}) = \overline{X} + \frac{\sigma\Phi^{-1}(1-\alpha/2)}{\sqrt{n}}$$

Example 2 : Using a Pivotal Quantity : Exponential Case

Suppose that $X_1, \ldots, X_n \sim \text{Exponential}(\theta)$. Then

$$T(\mathbf{X}) = \sum_{i=1}^{n} X_i \sim \text{Gamma}(n, \theta)$$

and hence

$$Q(\mathbf{X}, \theta) = \theta T(\mathbf{X}) \sim \text{Gamma}(n, 1)$$

is a pivotal quantity. We have that

$$\Pr[c_1 < Q(\mathbf{X}, \theta) < c_2; \theta] = 1 - \alpha$$

if c_1 and c_2 are the α_1 and α_2 quantiles of the $\text{Gamma}(n, 1)$ distribution. Hence a $1-\alpha$ interval is

$$[L(\mathbf{X}), U(\mathbf{X})] \equiv \left[\frac{c_1}{T(\mathbf{X})}, \frac{c_2}{T(\mathbf{X})}\right]$$

Example 3 : Using a Pivotal Quantity : Normal variance case
Suppose that $X_1, \ldots, X_{n_1} \sim \text{Normal}(\theta_1, \sigma_1^2)$ and $Y_1, \ldots, Y_{n_2} \sim \text{Normal}(\theta_2, \sigma_2^2)$ are independent random samples. Then

$$Q_1(\mathbf{X}, \sigma_1^2) = \frac{(n_1-1)s_1^2}{\sigma_1^2} = \frac{1}{\sigma_1^2}\sum_{i=1}^{n_1}(X_i - \overline{X})^2 \qquad Q_2(\mathbf{Y}, \sigma_2^2) = \frac{(n_2-1)s_2^2}{\sigma_2^2} = \frac{1}{\sigma_2^2}\sum_{i=1}^{n_2}(Y_i - \overline{Y})^2$$

are pivotal quantities, as

$$Q_1(\mathbf{X}, \sigma_1^2) \sim \text{Chisquared}(n_1 - 1) \qquad Q_2(\mathbf{X}, \sigma_2^2) \sim \text{Chisquared}(n_2 - 1).$$

To construct a $1-\alpha$ confidence interval, note that

$$\Pr[c_{11} < Q_1(\mathbf{X}, \sigma_1^2) < c_{12}; \sigma_1^2] = 1 - \alpha$$

if c_{11} and c_{12} are the α_1 and α_2 quantiles of the Chi-squared distribution with $n_1 - 1$ degrees of freedom. Hence with probability $1 - \alpha$

$$c_{11} < Q_1(\mathbf{X}, \sigma_1^2) < c_{12} \qquad \therefore \qquad c_{11} < \frac{(n_1-1)s_1^2}{\sigma_1^2} < c_{12}$$

and therefore the $1-\alpha$ interval is

$$[L_1(\mathbf{X}), U_1(\mathbf{X})] \equiv \left[\frac{(n_1-1)s_1^2}{c_{12}}, \frac{(n_1-1)s_1^2}{c_{11}}\right]$$

with a similar interval for σ_2^2. Note also that by previous results

$$\frac{Q_1(\mathbf{X}, \sigma_1^2)/(n_1-1)}{Q_2(\mathbf{X}, \sigma_2^2)/(n_2-1)} = \frac{s_1^2}{s_2^2}\frac{\sigma_2^2}{\sigma_1^2} \sim Fisher(n_1 - 1, n_2 - 1)$$

is also a pivotal quantity, so by similar arguments to the above, a $1-\alpha$ interval for σ_1^2/σ_2^2 is

$$\left[\frac{s_1^2}{s_2^2 c_2}, \frac{s_1^2}{s_2^2 c_1}\right]$$

where c_1 and c_2 are the $\alpha/2$ and $1-\alpha/2$ quantiles of the $Fisher(n_1 - 1, n_2 - 1)$ distribution.

Example 4 : Inverting a Likelihood Ratio Statistic : Exponential case
Suppose that $X_1, \ldots, X_n \sim \text{Exponential}(\theta)$ and we wish to test the hypotheses

$$\begin{aligned} H_0 &: \theta = \theta_0 \\ H_1 &: \theta \neq \theta_0 \end{aligned}$$

The likelihood ratio test for these hypotheses is based on the statistic

$$\lambda_{\mathbf{X}}(\mathbf{x}) = \frac{\sup_{\theta \in \Theta_0} f_{\mathbf{X}}(\mathbf{x}; \theta)}{\sup_{\theta \in \Theta} f_{\mathbf{X}}(\mathbf{x}; \theta)}.$$

Under H_1, the ML estimator of θ is $\widehat{\theta}_n = 1/\overline{X}$, so

$$\lambda_{\mathbf{X}}(\mathbf{X}) = \frac{\theta_0^n \exp\{-n\theta_0 \overline{X}\}}{\widehat{\theta}^n \exp\{-n\widehat{\theta}\overline{X}\}} = \left(\frac{\theta_0 T(\mathbf{X})}{n}\right)^n \exp\{-T(\mathbf{X})\theta_0 + n\}$$

where, for any θ,

$$T(\mathbf{X}) = \sum_{i=1}^{n} X_i \sim \text{Gamma}(n, \theta) \qquad \therefore \qquad \theta T(\mathbf{X}) \sim \text{Gamma}(n, 1).$$

The acceptance region $\mathcal{A}(\theta_0)$ is the set

$$\{\mathbf{x} : \lambda_{\mathbf{X}}(\mathbf{x}) \geq c_1\}$$

which is equivalent to the set

$$\{t : (\theta_0 t)^n \exp\{-t\theta_0\} \geq c_2\}.$$

In general, there are two solutions $a_1(\theta_0) < a_2(\theta_0)$ to the equation

$$(\theta_0 t)^n \exp\{-t\theta_0\} = c_2 \tag{1}$$

or equivalently

$$n \log t - \theta_0 t = c_3 \tag{2}$$

but the solutions can only be found numerically; we must choose c_3 such that

$$\Pr[a_1(\theta_0) < T(\mathbf{X}) < a_2(\theta_0); \theta_0] = 1 - \alpha. \tag{3}$$

In practice, we might choose a range of values of c_3, then find $a_1(\theta_0)$ and $a_2(\theta_0)$ as solutions to equation (2), and then check equation (3) to see whether the probability is matched. In Figure 1 below, the acceptance region is computed for $n = 10, \theta_0 = 5$ and $\alpha = 0.05$

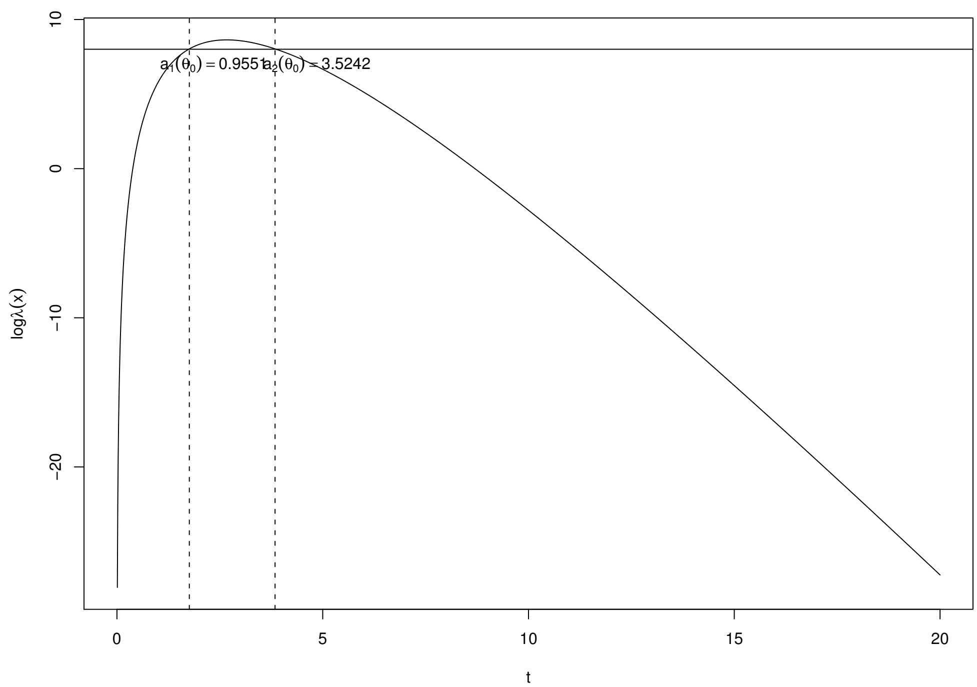

Figure 1: The $\alpha = 0.05$ acceptance region, $\mathcal{A}(\theta_0)$, for the Exponential model with $\theta_0 = 5$ and $n = 10$ is $(0.9551, 3.542)$. We move the value of k_3 up the y-axis until the intersection points, $a_1(\theta_0)$ and $a_2(\theta_0)$, of the horizontal line and the function $g(t) = n \log t - \theta_0 t$ define a region containing probability $1 - \alpha$.

To invert $\mathcal{A}(\theta_0)$ to get the $1 - \alpha$ confidence interval, we seek, for fixed data $\mathbf{x}$ and summary statistic $T(\mathbf{x})$, the set

$$\mathcal{C}(T(\mathbf{x})) = \{\theta : T(\mathbf{x}) \in \mathcal{A}(\theta)\} = \{\theta : (\theta T(\mathbf{x}))^n \exp\{-\theta T(\mathbf{x})\} \geq k_2\}$$

As the distribution is unimodal, a $1-\alpha$ confidence interval must take the form

$$\mathcal{C}(T(\mathbf{x})) = \{\theta : L(T(\mathbf{x})) \leq \theta \leq U(T(\mathbf{x}))\}$$

Writing $t = T(\mathbf{x})$, from equation (1) and by analogy with Figure 1, we must have

$$(tL(t))^n \exp\{-tL(t)\} = (tU(t))^n \exp\{-tU(t)\}. \tag{4}$$

If $a = tL(t)$ and $b = tU(t)$, then the interval is

$$\{\theta : a/t \leq \theta \leq b/t\}$$

where a and b satisfy

$$\Pr[a/T \leq \theta \leq b/T; \theta] = \Pr[a \leq \theta T \leq b] = 1-\alpha$$

where $\theta T \sim \text{Gamma}(n, 1)$. Thus from (4) we require that

$$a^n e^{-a} = b^n e^{-b}$$

whilst

$$\int_a^b \frac{1}{\Gamma(n)} x^{n-1} e^{-x}\, dx = 1-\alpha.$$

Therefore, solving for a and b is numerically straightforward using a look-up table approach. The code below in R demonstrates how this might be done; for a fine grid $\epsilon, 2\epsilon, \ldots, \alpha-\epsilon$, we compute the quantiles q_L and q_U corresponding to probabilities $m\epsilon$ and $m\epsilon + 1 - \alpha$, and then find the value of m such that

$$q_L^n e^{-q_L} - q_U^n e^{-q_U}$$

is as close as possible to zero.

```
n<-10
eps<-1e-6
eps.vec<-seq(eps,alpha-eps,by=eps)
qL.vec<-qgamma(eps.vec,n,1)
qU.vec<-qgamma(eps.vec+1-alpha,n,1)
d.vec<-exp(n*log(qL.vec)-qL.vec)-exp(n*log(qU.vec)-qU.vec)
a<-aL.vec[which.min(d.vec*d.vec)]
b<-qU.vec[which.min(d.vec*d.vec)]
```

which yields the following results

n	5	10	15	20	25	30	35	40	45	50
a	1.758	4.979	8.603	12.439	16.412	20.482	24.626	28.829	33.080	37.372
b	10.864	17.613	23.979	30.137	36.162	42.089	47.943	53.739	59.488	65.195

Note that this computation is independent of $t = T(\mathbf{x})$; to obtain the confidence interval, we need to divide a and b by t. For example, if $n = 10$ and $t = T(\mathbf{x}) = 2.281$, we have

$$L(T(\mathbf{x})) = \frac{4.979}{2.281} = 2.183 \qquad U(T(\mathbf{x})) = \frac{17.613}{2.281} = 7.722$$

Note that as the distribution of $Q(\mathbf{X}, \theta) = \theta T(\mathbf{X})$ does not depend on θ, it is a pivotal quantity, so

$$\Pr[a \leq \theta T \leq b] = \Pr[a/T \leq \theta \leq b/T] = 1-\alpha$$

already yields a $1-\alpha$ confidence interval; the additional constraint in equation (4) ensures that the interval is as short as possible.